

Rank-one Convexification for Sparse Regression

Alper Atamtürk

ATAMTURK@BERKELEY.EDU

*Department of Industrial Engineering & Operations Research
University of California
Berkeley, CA 94720, USA*

Andrés Gómez

GOMEZAND@USC.EDU

*Department of Industrial & Systems Engineering
University of Southern California
Los Angeles, CA 90089, USA*

Editor: Vahab Mirrokni

Abstract

Sparse regression models are increasingly prevalent due to their ease of interpretability and superior out-of-sample performance. However, the exact model of sparse regression with an ℓ_0 -constraint restricting the support of the estimators is a challenging (\mathcal{NP} -hard) non-convex optimization problem. In this paper, we derive new strong convex relaxations for sparse regression. These relaxations are based on the convex-hull formulations for rank-one quadratic terms with indicator variables. The new relaxations can be formulated as semidefinite optimization problems in an extended space and are stronger and more general than the state-of-the-art formulations, including the perspective reformulation and formulations with the reverse Huber penalty and the minimax concave penalty functions. Furthermore, the proposed rank-one strengthening can be interpreted as a *non-separable, non-convex, unbiased sparsity-inducing regularizer*, which dynamically adjusts its penalty according to the shape of the error function without inducing bias for the sparse solutions. In our computational experiments with benchmark datasets, the proposed conic formulations are solved within seconds and result in near-optimal solutions (with 0.4% optimality gap on average) for non-convex ℓ_0 -problems. Moreover, the resulting estimators also outperform alternative convex approaches, such as lasso and elastic net regression, from a statistical perspective, achieving high prediction accuracy and good interpretability.

Keywords: sparse regression, best subset selection, lasso, elastic net, conic formulations, non-convex regularization

1 Introduction

Given a model matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ of explanatory variables, a vector $\mathbf{y} \in \mathbb{R}^n$ of response variables, regularization parameters $\lambda, \mu \geq 0$ and a desired sparsity level $k \in \mathbb{Z}_+$, we consider the least squares regression problem

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2 + \mu\|\boldsymbol{\beta}\|_1 \text{ s.t. } \|\boldsymbol{\beta}\|_0 \leq k, \quad (1)$$

where $\|\boldsymbol{\beta}\|_0$ denotes cardinality of the support of $\boldsymbol{\beta}$. Problem (1) encompasses a broad range of regression models. It includes as special cases: **ridge** regression (Hoerl and Kennard, 1970), when $\lambda > 0$, $\mu = 0$ and $k \geq p$; **lasso** (Tibshirani, 1996), when $\lambda = 0$, $\mu \geq 0$ and $k \geq p$;

elastic net (Zou and Hastie, 2005) when $\lambda, \mu > 0$ and $k \geq p$; **best subset selection** (Miller, 2002), when $\lambda = \mu = 0$ and $k < p$. Additionally, Bertsimas and Van Parys (2020) propose to solve (1) with $\lambda > 0$, $\mu = 0$ and $k < p$ for high-dimensional regression problems, while Mazumder et al. (2022) study (1) with $\lambda = 0$, $\mu > 0$ and $k < p$ for problems with low Signal-to-Noise Ratios (SNR). The results in this paper cover all versions of (1) with $k < p$; moreover, they can be extended to problems with non-separable regularizations of the form $\lambda \|\mathbf{A}\boldsymbol{\beta}\|_2^2 + \mu \|\mathbf{C}\boldsymbol{\beta}\|_1$, resulting in sparse variants of the **fused lasso** (Padilla et al., 2017; Tibshirani et al., 2005), **generalized lasso** (Lin et al., 2014; Tibshirani and Taylor, 2011) and **smooth lasso** (Hebiri et al., 2011), among others.

1.1 Regularization Techniques

The motivation and benefits of regularization and sparsity are well-documented in the literature. In particular, generalization and interpretability are the key drivers of sparsity in machine learning models. Generalization refers to the ability of a model to perform well out-of-sample. The principle of parsimony postulates that observed phenomena often admit simple explanations, and thus sparse models are preferable as they are more likely to capture such explanations; in fact, Hastie et al. (2001) coined the bet on sparsity principle, i.e., using an inference procedure that performs well in sparse problems since no procedure can do well in dense problems. Interpretability, on the other hand, is becoming increasingly important due to the deployment of machine learning models in high-stakes situations (Rudin, 2019), where complex models can sometimes lead to results that are unfair, discriminatory, or otherwise undesirable. As such, interpretability is critical in healthcare (Ustun and Rudin, 2016, 2019; Wanjiru et al., 2021) and public policy (Aghaei et al., 2019; Azizi et al., 2018; Zeng et al., 2017) settings for example. Moreover, interpretable learning models are also preferable when the output of the model is intended to serve as input to a downstream decision-making process (Cozad et al., 2014; Lombardi et al., 2017; Maragno et al., 2021).

Best subset selection with $k < p$ and $\lambda = \mu = 0$ is the direct approach to enforce sparsity without introducing bias. In contrast, ridge regression with $\lambda > 0$ (Tikhonov regularization) is known to induce shrinkage and bias, which can be desirable, for example, when \mathbf{X} is not orthogonal, but it does not encourage sparsity. On the other hand, **lasso**, employs ℓ_1 regularization with $\mu > 0$ to achieve both shrinkage and sparsity. However, the inability to separately control for shrinkage and sparsity may result in subpar performance in some cases (Kelner et al., 2024; Miller, 2002; Zhang et al., 2008, 2012, 2014; Zhao and Yu, 2006; Zou, 2006). Moreover, achieving a target sparsity level k with **lasso** requires significant experimentation with the penalty parameter μ (Chichignoud et al., 2016). When $k \geq p$, the ℓ_0 cardinality constraint becomes redundant, and (1) reduces to a convex optimization problem and can be solved easily. On the other hand, when $k < p$, problem (1) is non-convex and \mathcal{NP} -hard (Natarajan, 1995), thus finding an optimal solution may require excessive computational effort and methods to solve it approximately are used instead (Huang et al., 2018; Nevo and Ritov, 2017), often relying on local heuristics (Hazimeh and Mazumder, 2020; Zhu et al., 2022, 2020). Due to the perceived difficulties of tackling the non-convex ℓ_0 constraint in (1), **lasso**-type simpler approaches continue to be preferred for inference problems with sparsity (Hastie et al., 2015), and have also been incorporated with

other learning models such as deep neural networks (Cherepanova et al., 2024; Dinh and Ho, 2020; Lemhadri et al., 2021).

Nonetheless, there has been a substantial effort to develop sparsity-inducing methodologies that do not incur as much shrinkage and bias as `lasso` does. These efforts have led to optimization problems of the form

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{i=1}^p \rho_i(\beta_i) \quad (2)$$

where $\rho_i : \mathbb{R} \rightarrow \mathbb{R}$ represent non-convex regularization functions. Examples of such regularization functions include ℓ_q penalties with $0 < q < 1$ (Frank and Friedman, 1993) and Smoothly Clipped Absolute Deviation (SCAD) penalty (Fan and Li, 2001). Although optimal solutions of (2) with non-convex regularizations may substantially improve upon the estimators obtained by `lasso`, solving (2) to optimality is still a difficult task (Hunter and Li, 2005; Mazumder et al., 2011; Zou and Li, 2008), and suboptimal solutions may not exhibit the improved statistical properties. To mitigate such difficulties, Zhang et al. (2010) propose the **minimax concave penalty** (MC_+), a class of sparsity-inducing penalty functions where the non-convexity of ρ is offset by the convexity of $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$ for sufficiently sparse solutions, so that (2) remains convex – Zhang et al. (2010) refer to this property as sparse convexity. Thus, in the ideal scenario (and with proper tuning of the parameter controlling the concavity of ρ), the MC_+ penalty is able to retain the sparsity and unbiasedness of **best subset selection** while preserving convexity, resulting in the best of both worlds. However, due to the *separable* form of the regularization term, the effectiveness of MC_+ greatly depends on the diagonal dominance of the matrix $\mathbf{X}^\top \mathbf{X}$ (this dependency will be discussed in more detail in §3), and may result in poor performance when the diagonal dominance is low.

Yet, in many practical applications, the matrix $\mathbf{X}^\top \mathbf{X}$ has low eigenvalues and lacks diagonal dominance. To illustrate, Table 1 presents the diagonal dominance of five datasets from the UCI Machine Learning Repository (Dheeru and Karra Taniskidou, 2017) used by Gómez and Prokopyev (2021) and Miyashiro and Takano (2015), as well as the **diabetes** dataset with all second interactions used by Bertsimas et al. (2016) and Efron et al. (2004). The diagonal dominance of a positive semidefinite matrix \mathbf{A} is computed as

$$\text{dd}(\mathbf{A}) := (1/\text{tr}(\mathbf{A})) \max_{\mathbf{d} \in \mathbb{R}_+^p} \mathbf{e}^\top \mathbf{d} \text{ s.t. } \mathbf{A} - \text{diag}(\mathbf{d}) \succeq 0,$$

where \mathbf{e} is the p -dimensional vector of ones, $\text{diag}(\mathbf{d})$ is the diagonal matrix such that $\text{diag}(\mathbf{d})_{ii} = d_i$ and $\text{tr}(\mathbf{A})$ denotes the trace of \mathbf{A} . Accordingly, the diagonal dominance is the trace of the largest diagonal matrix that can be extracted from \mathbf{A} without violating positive semidefiniteness, divided by the trace of \mathbf{A} . Table 1 shows that the diagonal dominance of $\mathbf{X}^\top \mathbf{X}$ is low or even 0% in some instances, and MC_+ struggles for these datasets as we demonstrate in §5.

1.2 Mixed-Integer Optimization Formulations

An alternative to using nonconvex regularizers is to leverage recent advances in mixed-integer optimization (MIO) to tackle (1) directly (Bertsimas and King, 2015; Bertsimas

dataset	p	n	dd×100%
housing	13	506	26.7%
servo	19	167	0.0%
auto MPG	25	392	1.5%
solar flare	26	1,066	8.8%
breast cancer	37	196	3.6%
diabetes	64	442	0.0%
crime	100	1993	13.5 %

 Table 1: Diagonal dominance of $\mathbf{X}^\top \mathbf{X}$ for benchmark datasets.

et al., 2016; Cozad et al., 2014). By introducing binary indicator variables $\mathbf{z} \in \{0, 1\}^p$, where $z_i = \mathbb{1}_{\beta_i \neq 0}$, problem (1) can be reformulated as

$$\mathbf{y}^\top \mathbf{y} + \min_{\beta, \mathbf{z}, \mathbf{u}} -2\mathbf{y}^\top \mathbf{X}\beta + \beta^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \beta + \mu \sum_{i=1}^p u_i \quad (3a)$$

$$\text{s.t. } \sum_{i=1}^p z_i \leq k \quad (3b)$$

$$\beta_i \leq u_i, \quad -\beta_i \leq u_i \quad i = 1, \dots, p \quad (3c)$$

$$\beta_i(1 - z_i) = 0 \quad i = 1, \dots, p \quad (3d)$$

$$\beta \in \mathbb{R}^p, \quad \mathbf{z} \in \{0, 1\}^p, \quad \mathbf{u} \in \mathbb{R}_+^p. \quad (3e)$$

The nonconvexity of (1) is captured by the complementary constraints (3d) and the integrality constraints $\mathbf{z} \in \{0, 1\}^p$. In fact, one of the main challenges to solve (3) is handling constraints (3d). A standard approach in the MIO literature is to replace (3d) with the so-called big- M constraints

$$-Mz_i \leq \beta_i \leq Mz_i \quad (4)$$

for a sufficiently large number M to bound the variables β_i . However, these big- M constraints (4) are poor approximations of constraints (3d), *especially in the case of regression problems where no natural big- M value is available*. Bertsimas et al. (2016) propose approaches to compute provable big- M values, but such values often result in prohibitively large computational times even in problems with a few dozens variables (or, even worse, may lead to numerical instabilities and cause convex solvers to crash). Alternatively, heuristic values for the big- M values can be estimated, e.g., by setting $M = \tau \|\hat{\beta}\|_\infty$, where $\tau \in \mathbb{R}_+$ and $\hat{\beta}$ is a feasible solution of (1) found via a heuristic method¹. While using such heuristic values yield reasonable performance for small enough values of τ , they may eliminate optimal solutions. Due to the modeling power of MIO, formulation (3) can be further extended to incorporate additional considerations such as outliers (Insolia et al., 2022), or graphical structures arising when learning Bayesian networks (Küçükyavuz et al., 2023).

1. This method with $\tau = 2$ was used in the computations of Bertsimas et al. (2016).

Branch-and-bound algorithms for MIO leverage strong convex relaxations of problems to prune the search space and reduce the number of sub-problems to enumerate (and, in some cases, eliminate the need for enumeration altogether). Additionally, strong relaxations can also be used to design exact screening rules that quickly identify critical explanatory variables and provably eliminate irrelevant ones to reduce the dimension of the inference problems to solve (Atamtürk and Gómez, 2020; Deza and Atamtürk, 2022). Thus, a critical step to speed up the solution times for (3) is to derive convex relaxations that approximate the non-convex problem well (Atamtürk and Narayanan, 2007). Such strong relaxations can also be used directly to find good estimators for the inference problems (without branch-and-bound); in fact, it is well known that the natural convex relaxation of (3) with $\lambda = \mu = 0$ and big- M constraints is precisely **lasso**, see (Dong et al., 2015) for example. Therefore, sparsity-inducing techniques that more accurately capture the properties of the non-convex constraint $\|\beta\|_0 \leq k$ can be found by deriving tighter convex relaxations of (1). Pilanci et al. (2015) exploit the Tikhonov regularization term and convex analysis to construct an improved convex relaxation using the **reverse Huber penalty**. In a similar vein, Bertsimas and Van Parys (2020) leverage Tikhonov regularization and duality to give an efficient algorithm for high-dimensional sparse regression. We refer the reader to the recent paper by Tillmann et al. (2024) for a survey of solution approaches to problems with sparsity constraints, such as (3).

1.3 The Perspective Relaxation

Problem (3) is a mixed-integer convex quadratic optimization problem with indicator variables, a class of problems which has received a fair amount of attention in the optimization literature. In particular, the **perspective relaxation** (Aktürk et al., 2009; Frangioni and Gentile, 2006; Günlük and Linderoth, 2010) is, by now, a standard technique that can be used to substantially strengthen the convex relaxations by exploiting *separable* quadratic terms. Specifically, consider the mixed-integer epigraph of a one-dimensional quadratic function with an indicator constraint,

$$Q_1 = \{z \in \{0, 1\}, \beta \in \mathbb{R}, t \in \mathbb{R}_+ : \beta_i^2 \leq t, \beta_i(1 - z_i) = 0\}.$$

The convex hull of Q_1 is obtained by relaxing the integrality constraint to bound constraints and using the closure of the perspective function² of β_i^2 , expressed as a rotated cone constraint:

$$\text{cl conv}(Q_1) = \left\{ z \in [0, 1], \beta \in \mathbb{R}, t \in \mathbb{R}_+ : \frac{\beta_i^2}{z_i} \leq t \right\}.$$

2. We use the convention that $\frac{\beta_i^2}{z_i} = 0$ when $\beta_i = z_i = 0$ and $\frac{\beta_i^2}{z_i} = \infty$ if $z_i = 0$ and $\beta_i \neq 0$.

Xie and Deng (2020) apply the **perspective relaxation** to the separable quadratic regularization term $\lambda\|\boldsymbol{\beta}\|_2^2$, i.e., reformulate (3) as

$$\mathbf{y}^\top \mathbf{y} + \min_{\boldsymbol{\beta}, \mathbf{z}, \mathbf{u}} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\beta}^\top (\mathbf{X}^\top \mathbf{X}) \boldsymbol{\beta} + \lambda \sum_{i=1}^p \frac{\beta_i^2}{z_i} + \mu \sum_{i=1}^p u_i \quad (5a)$$

$$\text{s.t. } \sum_{i=1}^p z_i \leq k \quad (5b)$$

$$\beta_i \leq u_i, \quad -\beta_i \leq u_i, \quad i = 1, \dots, p \quad (5c)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \quad \mathbf{z} \in \{0, 1\}^p, \quad \mathbf{u} \in \mathbb{R}_+^p. \quad (5d)$$

Moreover, they show that the continuous relaxation of (5) is equivalent to the continuous relaxation of the formulation used by Bertsimas and Van Parys (2020). Dong et al. (2015) also study the **perspective relaxation** in the context of regression: first, they show that using the **reverse Huber penalty** (Pilanci et al., 2015) is, in fact, equivalent to just solving the convex relaxation of (5) — thus the relaxations of (Bertsimas and Van Parys, 2020; Pilanci et al., 2015; Xie and Deng, 2020) all coincide; second, they propose to use an *optimal perspective relaxation*, i.e., by applying the perspective relaxation to a separable quadratic function $\boldsymbol{\beta}^\top \mathbf{D}\boldsymbol{\beta}$, where \mathbf{D} is a nonnegative diagonal matrix such that $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} - \mathbf{D} \succeq 0$; finally, they show that solving this stronger convex relaxation of the *optimal perspective relaxation* is, in fact, equivalent to using the MC_+ penalty (Zhang et al., 2010). However, the authors also point out that if $\lambda = 0$ and a suitable matrix \mathbf{D} cannot be found, then the optimal perspective relaxation reduces to **lasso**. For example, from Table 1, we see that the optimal perspective relaxation would reduce to **lasso** in the **servo** and **diabetes** datasets.

The perspective relaxation is now a state-of-the-art method to convexify problems with separable terms and indicator variables. Recent works have extended perspective relaxations to more general classes of problems, such as inference problems with grouped variable selection constraints (Hazimeh et al., 2023b). However, there are relatively few convexification techniques for problems without separable terms (Frangioni et al., 2020; Gómez, 2021; Han et al., 2023; Jeon et al., 2017). In fact, among the previously discussed methods for sparse regression, the *optimal perspective relaxation* of Dong et al. (2015) is the only one that does not explicitly require the use of the Tikhonov regularization $\lambda\|\boldsymbol{\beta}\|_2^2$. Nonetheless, as the authors point out, if $\lambda = 0$ then the method is effective only when the matrix $\mathbf{X}^\top \mathbf{X}$ is sufficiently diagonally dominant, which, as illustrated in Table 1, is not necessarily the case in practice. As a consequence, **perspective relaxation** techniques may be insufficient to tackle problems when large shrinkage is undesirable and, λ is small.

1.4 Contributions

In this paper we derive stronger convex relaxations of (3) than the *optimal perspective relaxation*. These relaxations are obtained from the study of ideal (convex-hull) formulations of the mixed-integer epigraphs of *non-separable rank-one quadratic functions with indicators*. Since the **perspective relaxation** corresponds to the ideal formulation of a *one-dimensional* rank-one quadratic function, the proposed relaxations generalize and

strengthen the existing results. In particular, they *dominate perspective relaxation* approaches for all values of the regularization parameter λ and, critically, are able to achieve high-quality approximations of (1) even in low diagonal dominance settings with $\lambda = 0$. Alternatively, our results can also be interpreted as a new *non-separable, non-convex, unbiased* regularization penalty $\rho_{\text{R1}}(\boldsymbol{\beta})$ which: (i) imposes larger penalties than the separable minimax concave penalty (Zhang et al., 2010) $\rho_{\text{MC}_+}(\boldsymbol{\beta})$ to dense estimators, thus achieving better sparsity-inducing properties; and (ii) the nonconvexity of the penalty function is offset by the convexity of the term $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$, and the resulting continuous problem can be solved to global optimality using convex optimization tools. In fact, they can be formulated as semidefinite optimization and, in certain special cases, as conic quadratic optimization. We point out that in recent papers, some of the relaxations proposed in this paper have been extended to more general settings with non-quadratic functions, low rank (but not necessarily rank-one) functions, and side constraints (Han and Gómez, 2024; Shafiee and Kılınç-Karzan, 2024).

To illustrate the regularization point of view for the proposed relaxations, consider a two-predictor regression problem in Lagrangian form:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 + \mu \|\boldsymbol{\beta}\|_1 + \kappa \|\boldsymbol{\beta}\|_0, \quad (6)$$

where $\mathbf{X}^\top \mathbf{X} = \begin{pmatrix} 1 + \delta & 1 \\ 1 & 1 + \delta \end{pmatrix}$ and $\delta \geq 0$ is a parameter controlling the diagonal dominance. Figure 1 depicts the graphs of well-known regularizations including **lasso** ($\lambda = \kappa = 0$, $\mu = 1$), **ridge** ($\mu = \kappa = 0$, $\lambda = 1$), **elastic net** ($\kappa = 0$, $\lambda = \mu = 0.5$), the MC_+ penalty for different values of δ and the proposed rank-one **R1** regularization. The graphs of MC_+ and **R1** are obtained by setting $\lambda = \mu = 0$ and $\kappa = 1$, and using the appropriate convex strengthening, see §3 for details. Observe that the **R1** regularization results in larger penalties than MC_+ for all values of δ , and the improvement increases as $\delta \rightarrow 0$. In addition, Figure 2 shows the effect of using the **lasso** constraint $\|\boldsymbol{\beta}\|_1 \leq k$, the MC_+ constraint $\rho_{\text{MC}_+}(\boldsymbol{\beta}) \leq k$, and the rank-one constraint $\rho_{\text{R1}}(\boldsymbol{\beta}) \leq k$ in a two-dimensional problem to achieve sparse solutions satisfying $\|\boldsymbol{\beta}\|_0 \leq 1$. Specifically, let

$$\varepsilon^* = \min_{\|\boldsymbol{\beta}\|_0 \leq 1} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$$

be the minimum residual error of a sparse solution of the least squares problem. Figure 2 shows in gray the (possibly dense) points satisfying $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \leq \varepsilon^*$, and it shows in color the set of feasible points satisfying $\rho(\boldsymbol{\beta}) \leq k$, where ρ is a given regularization and k is chosen so that the feasible region (color) intersects the level sets (gray). We see that neither **lasso** nor MC_+ is able to exactly recover an optimal sparse solution for any diagonal dominance parameter δ , despite significant shrinkage ($k < 1$). In contrast, the rank-one constraint $\rho_{\text{R1}}(\boldsymbol{\beta}) \leq k$ adapts to the curvature of the error function $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$ to induce higher sparsity: in particular, the “natural” constraint $\rho_{\text{R1}}(\boldsymbol{\beta}) \leq 1$, with the target sparsity $k = 1$, results in exact recovery without shrinkage in all cases.

Finally, Figure 3 shows the strength of relaxations of (1) discussed in this paper. The “big- M ” relaxation is the natural convex relaxation of (3) obtained by replacing $z \in \{0, 1\}^p$ by $z \in [0, 1]^p$, used by Bertsimas et al. (2016); Cozad et al. (2014). The perspective

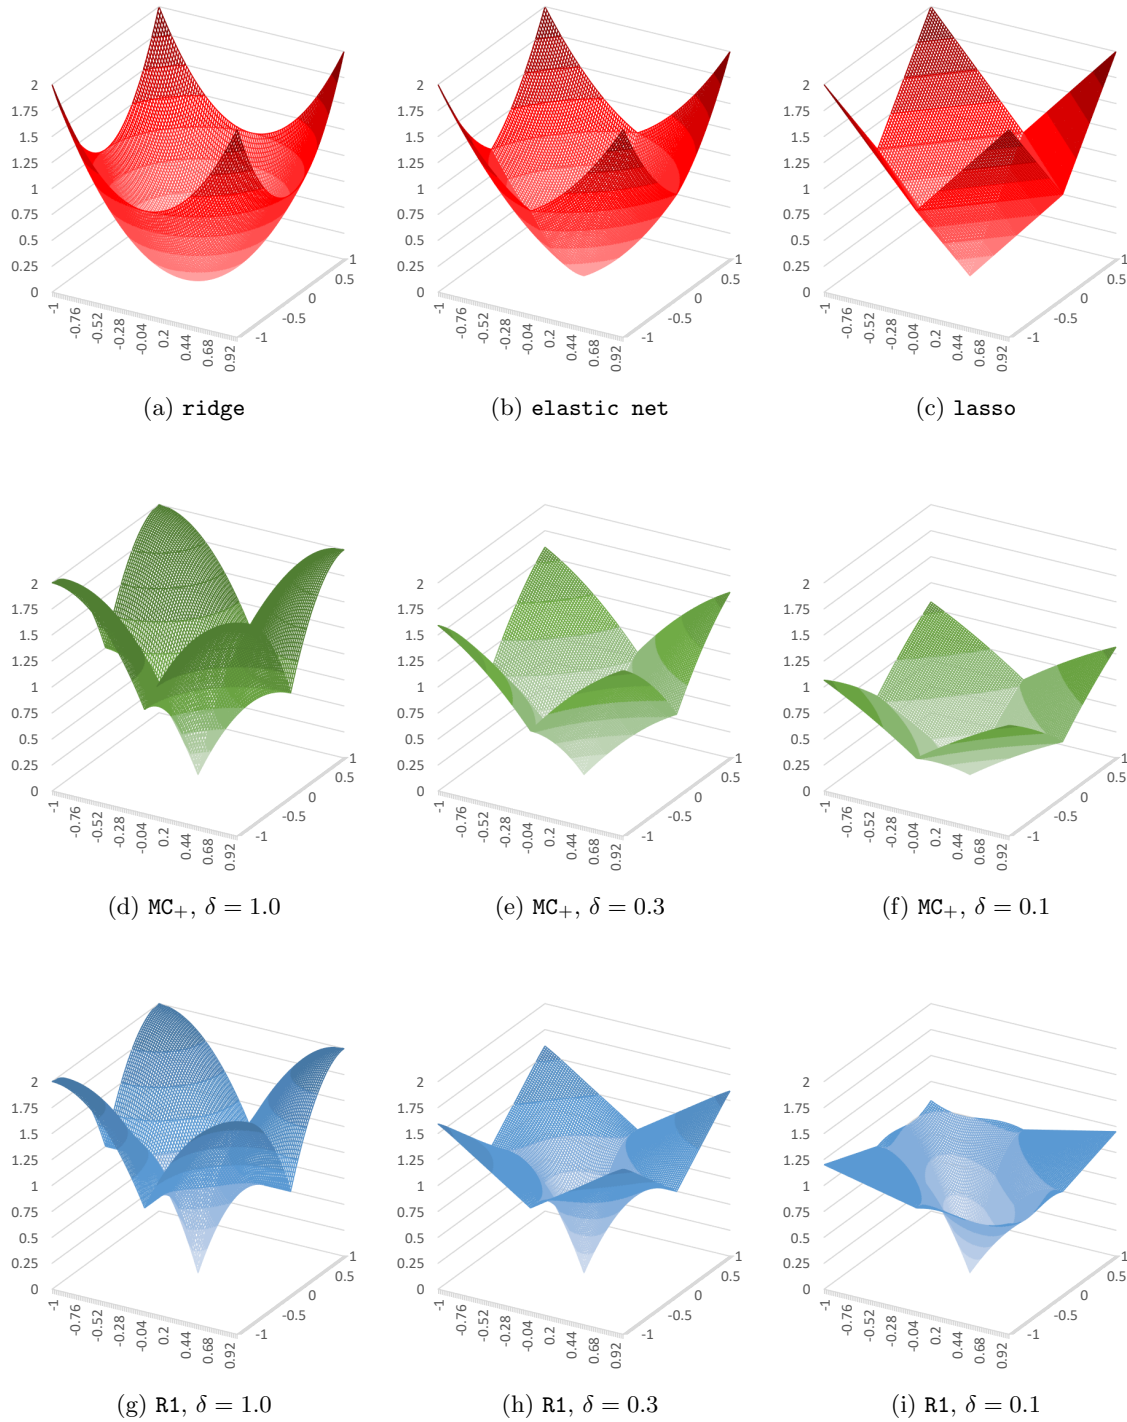


Figure 1: Graphs of regularization penalties with $p = 2$. The horizontal axes correspond to values of β_1 and β_2 , and the vertical axis corresponds to the regularization penalty. The **ridge**, **elastic net**, and **lasso** (top row) regularizations do not depend on the diagonal dominance, but induce substantial bias. The MC_+ regularization (second row) does not induce as much bias, but it depends on the diagonal dominance (δ). The new non-separable, non-convex $R1$ regularization (bottom row) induces larger penalties than MC_+ for all diagonal dominance values and is a closer approximation for the exact ℓ_0 penalty.

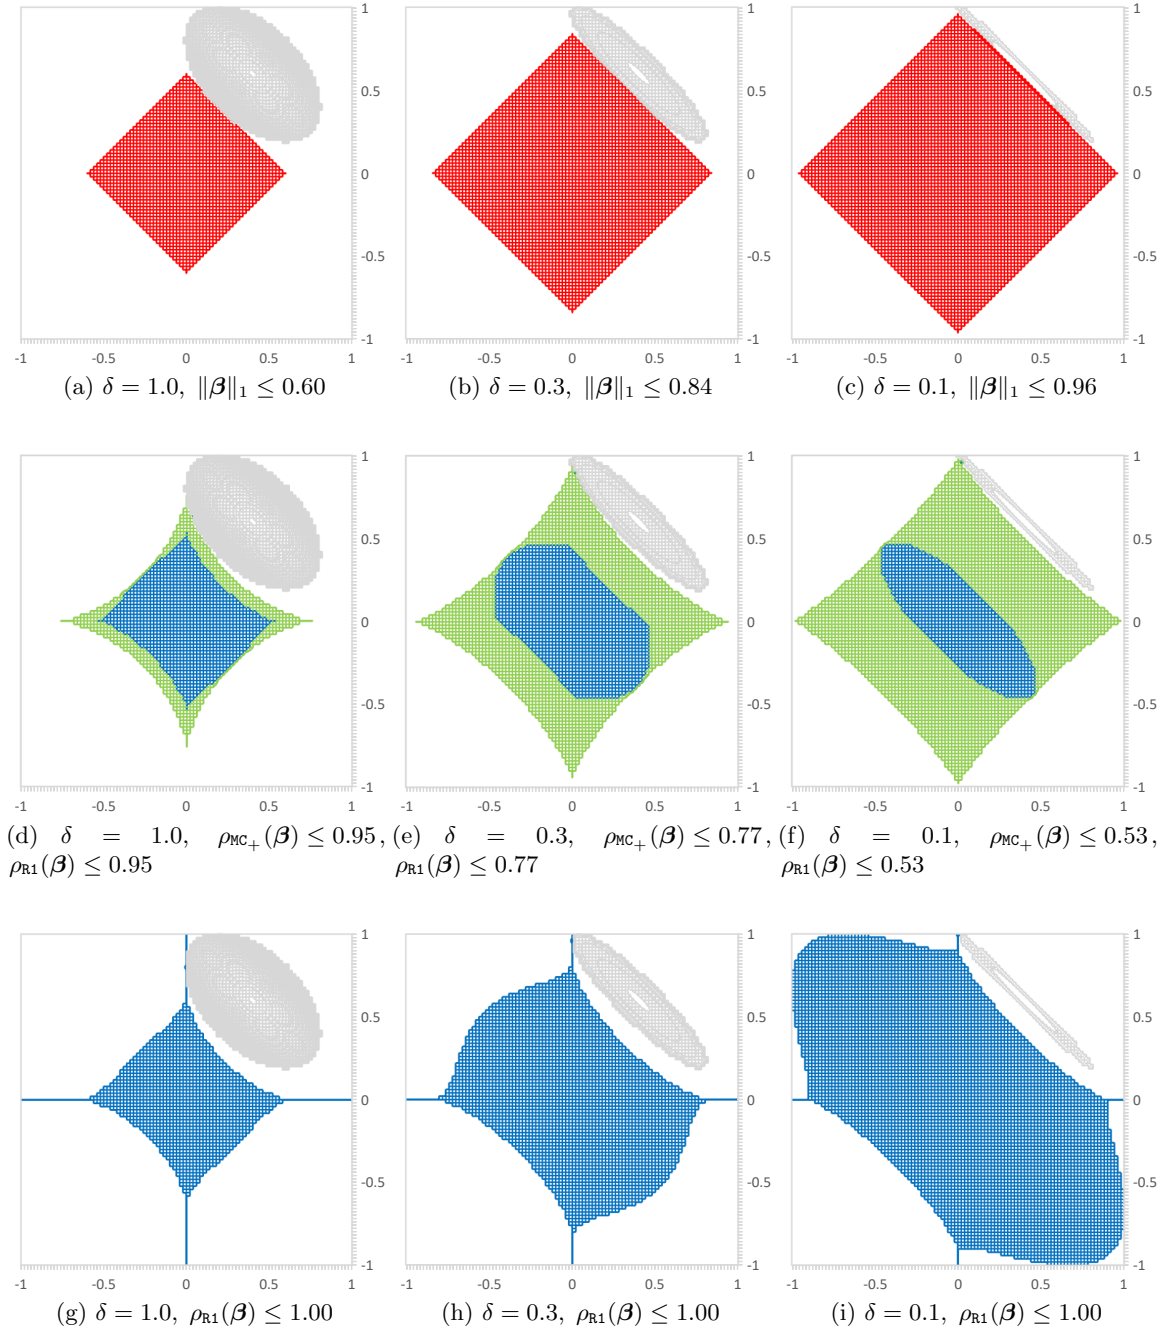


Figure 2: The axes correspond to the sparse solutions satisfying $\|\beta\|_0 \leq 1$. In gray: level sets given by $\|\mathbf{y} - \mathbf{X}\beta\|_2^2 \leq \varepsilon^*$; in red: feasible region for $\|\beta\|_1 \leq k$; in green: feasible region for $\rho_{\text{MC}_+}(\beta) \leq k$; in blue: feasible region for $\rho_{\text{R1}}(\beta) \leq k$. All lasso and MC_+ solutions above are dense even with significant shrinkage ($k < 1$). Rank-one constraint attains sparse solutions on the axes with no shrinkage ($k = 1$) for all diagonal dominance values δ .

relaxation is the natural convex relaxation of (5), which is the basis of recent methods (Bertsimas and Van Parys, 2020; Hazimeh et al., 2022; Pilanci et al., 2015; Xie and Deng, 2020; Hazimeh et al., 2023a) – note that this formulation may only be used if $\lambda > 0$. The “optimal perspective” relaxation, also referred to as sdp_1 in this paper, was explicitly given by Dong et al. (2015). Interestingly, it was recently shown (Han et al., 2022) that sdp_1 is equivalent to the standard Shor’s SDP relaxation (Shor, 1987) for problem (3)– a convex relaxation that has been proven to be very effective at approximating discrete optimization problems (Goemans and Williamson, 1995). This paper proposes new relaxations sdp_r , discussed in §2, which dominate all existing relaxations in terms of strength. It also proposes the new formulation sdp_{LB} , discussed in §4, which is easier to solve than sdp_r but still compares favorably with the “big- M ” and perspective formulations.

Outline The rest of the paper is organized as follows. In §2 we derive convex relaxations based on ideal formulations for rank-one quadratic terms with indicator variables. We also give an interpretation of the convex relaxations as unbiased regularizers and propose an explicit semidefinite optimization (SDP) formulation in an extended space, which can be implemented with off-the-shelf conic optimization solvers. In §3 we derive an explicit form of the regularization penalty for the two-dimensional case. In §4 we discuss the implementation of the proposed relaxation in a conic quadratic framework. In §5 we present computational experiments with synthetic as well as benchmark datasets, demonstrating that (i) the proposed formulation delivers near-optimal solutions (with provable optimality gaps) of (1) in most cases, (ii) using the proposed convex relaxation results in superior

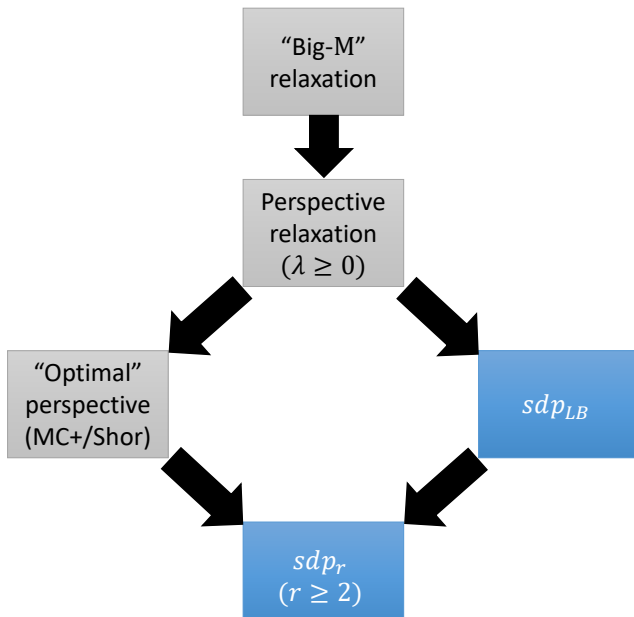


Figure 3: Strength of relaxations discussed in the paper. “ $A \Rightarrow B$ ” indicates that B is a stronger relaxation than A , i.e., is a better approximation for the non-convex problem (1). Blue boxes correspond to the new formulations proposed in this paper.

statistical performance when compared with the usual convex estimators such as elastic net. In §6, we conclude the paper with a few final remarks.

Notation Define $P = \{1, \dots, p\}$ and $\mathbf{e} \in \mathbb{R}^p$ be the vector of ones. Given $T \subseteq P$ and a vector $\mathbf{a} \in \mathbb{R}^p$, define \mathbf{a}_T as the subvector of \mathbf{a} induced by T , $a_i = \mathbf{a}_{\{i\}}$ as the i -th element of \mathbf{a} , and define $a(T) = \sum_{i \in T} a_i$. Given a symmetric matrix $\mathbf{A} \in \mathbb{R}^{p \times p}$, let \mathbf{A}_T be the submatrix of \mathbf{A} induced by $T \subseteq P$, and let \mathcal{S}_+^T be the set of $|T| \times |T|$ symmetric positive semidefinite matrices, i.e., $\mathbf{A}_T \succeq 0 \Leftrightarrow \mathbf{A}_T \in \mathcal{S}_+^T$. We use \mathbf{a}_T or \mathbf{A}_T to make explicit that a given vector or matrix is indexed by the elements of T or $T \times T$, respectively. Given matrices \mathbf{A}, \mathbf{B} of the same dimension, $\mathbf{A} \circ \mathbf{B}$ denotes the Hadamard product of \mathbf{A} and \mathbf{B} , and $\langle \mathbf{A}, \mathbf{B} \rangle$ denotes their inner product. Given a vector $\mathbf{a} \in \mathbb{R}^n$, let $\text{diag}(\mathbf{a})$ be the $n \times n$ diagonal matrix \mathbf{A} with $A_{ii} = a_i$. For a set $X \subseteq \mathbb{R}^p$, $\text{cl conv}(X)$ denotes the closure of the convex hull of X . Throughout the paper, we adopt the following convention for division by 0: given a scalar $s \geq 0$, $s/0 = \infty$ if $s > 0$ and $s/0$ if $s = 0$. For a scalar $a \in \mathbb{R}$, let $\text{sign}(a) = a/|a|$.

2 Convexification

In this section we introduce the proposed relaxations of problem (1). First, in §2.1, we describe the *ideal* relaxations for the mixed-integer epigraph of a rank-one quadratic term. Then, in §2.2, we use the relaxations derived in §2.1 to give strong relaxations of (1). Next, in §2.3, we give an interpretation of the proposed relaxations as unbiased sparsity-inducing regularizations. In §2.4 we present an explicit SDP representation of the proposed relaxations in an extended space. Finally, in §2.5, we comment on the strength of the proposed relaxations.

2.1 Rank-one Case

We first give a valid inequality for the mixed-integer epigraph of a convex quadratic function defined over the subsets of P . Given $\mathbf{A}_T \in \mathcal{S}_+^T$, consider the set

$$Q_T = \left\{ (\mathbf{z}, \boldsymbol{\beta}, t) \in \{0, 1\}^{|T|} \times \mathbb{R}^{|T|} \times \mathbb{R}_+ : \boldsymbol{\beta}^\top \mathbf{A}_T \boldsymbol{\beta} \leq t, \beta_i(1 - z_i) = 0, \forall i \in T \right\}.$$

Proposition 1 *The inequality*

$$\frac{\boldsymbol{\beta}^\top \mathbf{A}_T \boldsymbol{\beta}}{z(T)} \leq t \tag{7}$$

is valid for Q_T .

Proof Let $(\mathbf{z}, \boldsymbol{\beta}, t) \in Q_T$, and we verify that inequality (7) is satisfied. First observe that if $\mathbf{z} = \mathbf{0}$, then $\boldsymbol{\beta} = \mathbf{0}$ and inequality (7) reduces to $0 \leq t$, which is satisfied. Otherwise, if $z_i = 1$ for some $i \in T$, then $z(T) \geq 1$ and we find that $\frac{\boldsymbol{\beta}^\top \mathbf{A}_T \boldsymbol{\beta}}{z(T)} \leq \boldsymbol{\beta}^\top \mathbf{A}_T \boldsymbol{\beta} \leq t$, and inequality (7) is satisfied again. \blacksquare

Observe that if T is a singleton, i.e., $T = \{i\}$, then (7) reduces to the well-known perspective inequality $A_{ii}\beta_i^2 \leq tz_i$. Moreover, if $T = \{i, j\}$ and \mathbf{A}_T is rank-one, i.e., $\mathbf{A}_T = a_T a_T^\top$ with

$a_T = (a_i \ a_j)^\top$ and $\boldsymbol{\beta}^\top \mathbf{A}_T \boldsymbol{\beta} = |A_{ij}| \left(a\beta_i^2 \pm 2\beta_i\beta_j + (1/a)\beta_j^2 \right)$ for $A_{ij} = a_i a_j$ and $a = a_i/a_j$, then (7) reduces to

$$|A_{ij}| \left(a\beta_i^2 \pm 2\beta_i\beta_j + (1/a)\beta_j^2 \right) \leq t(z_i + z_j), \quad (8)$$

one of the inequalities proposed by Jeon et al. (2017) in the context of quadratic optimization with indicators and bounded continuous variables. Note that inequality (8) is, in general, weak for bounded continuous variables as non-negativity or other bounds can be used to strengthen the inequalities (Atamtürk and Gómez, 2018); and inequality (7) is, in general, weak for arbitrary matrices $\mathbf{A}_T \in \mathcal{S}_+^T$. Nonetheless, as we show next, inequality (7) is sufficient to describe the *ideal* (convex hull) description for Q_T if $\mathbf{A}_T = \mathbf{a}_T \mathbf{a}_T^\top$ is a rank-one matrix. Consider the special case of Q_T defined with a rank-one matrix:

$$Q_T^{r1} = \left\{ (\mathbf{z}, \boldsymbol{\beta}, t) \in \{0, 1\}^{|T|} \times \mathbb{R}^{|T|} \times \mathbb{R}_+ : (\mathbf{a}_T^\top \boldsymbol{\beta})^2 \leq t, \beta_i(1 - z_i) = 0, \forall i \in T \right\}.$$

Theorem 2 *If $a_i \neq 0$ for all $i \in T$, then*

$$cl \ conv(Q_T^{r1}) = \left\{ (\mathbf{z}, \boldsymbol{\beta}, t) \in [0, 1]^{|T|} \times \mathbb{R}^{|T|} \times \mathbb{R}_+ : (\mathbf{a}_T^\top \boldsymbol{\beta})^2 \leq t, \frac{(\mathbf{a}_T^\top \boldsymbol{\beta})^2}{z(T)} \leq t \right\}.$$

Proof Consider the optimization of an arbitrary linear function over Q_T^{r1} and $\bar{Q}_T := \left\{ (\mathbf{z}, \boldsymbol{\beta}, t) \in [0, 1]^{|T|} \times \mathbb{R}^{|T|} \times \mathbb{R}_+ : (\mathbf{a}_T^\top \boldsymbol{\beta})^2 \leq t, \frac{(\mathbf{a}_T^\top \boldsymbol{\beta})^2}{z(T)} \leq t \right\}$:

$$\min_{(\mathbf{z}, \boldsymbol{\beta}, t) \in Q_T^{r1}} \mathbf{u}_T^\top \mathbf{z} + \mathbf{v}_T^\top \boldsymbol{\beta} + \kappa t, \quad (9)$$

$$\min_{(\mathbf{z}, \boldsymbol{\beta}, t) \in \bar{Q}_T} \mathbf{u}_T^\top \mathbf{z} + \mathbf{v}_T^\top \boldsymbol{\beta} + \kappa t, \quad (10)$$

where $\mathbf{u}_T, \mathbf{v}_T \in \mathbb{R}^{|T|}$ and $\kappa \in \mathbb{R}$. We now show that either there exists an optimal solution of (10) that is feasible for (9), hence also optimal for (9) as \bar{Q}_T is a relaxation of Q_T^{r1} , or that (9) and (10) are both unbounded.

Observe that if $\kappa < 0$, then letting $\mathbf{z} = \boldsymbol{\beta} = \mathbf{0}$ and $t \rightarrow \infty$ we see that both problems are unbounded. If $\kappa = 0$ and $\mathbf{v}_T = \mathbf{0}$, then (10) reduces to $\min_{\mathbf{z} \in [0, 1]^{|T|}} \mathbf{u}_T^\top \mathbf{z}$, which has an optimal integral solution \mathbf{z}^* , and $(\mathbf{z}^*, \mathbf{0}, 0)$ is optimal for (9) and (10). If $\kappa = 0$ and $v_i \neq 0$ for some $i \in T$, then letting $\beta_i \rightarrow \pm\infty$, $z_i = 1$, and $\beta_j = z_j = t = 0$ for $j \neq i$, we find that both problems are unbounded. Thus, we may assume, without loss of generality that $\kappa > 0$, and, by scaling, $\kappa = 1$.

Additionally, as \mathbf{a}_T has no zero entry, we may assume, without loss of generality, that $\mathbf{a}_T = \mathbf{e}_T$, since otherwise $\boldsymbol{\beta}$ and \mathbf{v}_T can be scaled by letting $\bar{\beta}_i = a_i \beta_i$ and $\bar{v}_i = v_i/a_i$ to arrive at an equivalent problem. Moreover, a necessary condition for (9)–(10) to be bounded is that

$$-\infty < \min_{\boldsymbol{\beta} \in \mathbb{R}^{|T|}} \mathbf{v}_T^\top \boldsymbol{\beta} \text{ s.t. } \beta(T) = \zeta \quad (11)$$

for any fixed $\zeta \in \mathbb{R}$. It is easily seen that (11) has an optimal solution if and only if $v_i = v_j$ for all $i \neq j$. Thus, we may also assume without loss of generality that $\mathbf{v}_T^\top \boldsymbol{\beta} = v_0 \beta(T)$ for some scalar v_0 . Performing the above simplifications, we find that (10) reduces to

$$\min_{\mathbf{z} \in [0, 1]^{|T|}, \boldsymbol{\beta} \in \mathbb{R}^{|T|}, t \in \mathbb{R}} \mathbf{u}_T^\top \mathbf{z} + v_0 \beta(T) + t \text{ s.t. } \beta(T)^2 \leq t, \beta(T)^2 \leq tz(T). \quad (12)$$

Since the one-dimensional optimization $\min_{\beta \in \mathbb{R}} \{v_0\beta + \beta^2\}$ has an optimal solution, it follows that (12) is bounded and has an optimal solution. We now prove that (12) has an optimal solution that is integral in \mathbf{z} and satisfies $\boldsymbol{\beta} \circ (\mathbf{e} - \mathbf{z}) = 0$.

Let $(\mathbf{z}^*, \boldsymbol{\beta}^*, t^*)$ be an optimal solution of (12). First note that if $0 < z^*(T) < 1$, then $(\gamma\mathbf{z}^*, \gamma\boldsymbol{\beta}^*, \gamma t^*)$ is feasible for (10) for γ sufficiently close to 1, with objective value $\gamma(\mathbf{u}_T^\top \mathbf{z}^* + v_0\beta^*(T) + t^*)$. If $\mathbf{u}_T^\top \mathbf{z}^* + v_0\beta^*(T) + t^* \geq 0$, then for $\gamma = 0$, $(\gamma\mathbf{z}^*, \gamma\boldsymbol{\beta}^*, \gamma t^*)$ has an objective value equal or lower. Otherwise, for $\gamma = 1/z^*(T)$, $(\gamma\mathbf{z}^*, \gamma\boldsymbol{\beta}^*, \gamma t^*)$ is feasible and has a lower objective value. Thus, we find that either $\mathbf{0}$ is optimal for (12) (and the proof is complete), or there exists an optimal solution with $z^*(T) \geq 1$. In the later case, observe that any $(\bar{\mathbf{z}}, \bar{\boldsymbol{\beta}}, t^*)$ with $\bar{\mathbf{z}} \in \arg \min\{\mathbf{u}_T^\top \mathbf{z} : z^*(T) \geq 1, z \in [0, 1]^{|T|}\}$ is also optimal for (12), in particular there exists an optimal solution with $\bar{\mathbf{z}}$ integral.

Finally, let $i \in T$ be any index with $\bar{z}_i = 1$. Setting $\bar{\beta}_i = \beta^*(T)$ and $\bar{\beta}_j = 0$ for $i \neq j$, we find another optimal solution $(\bar{\mathbf{z}}, \bar{\boldsymbol{\beta}}, t^*)$ for (12) that satisfies the complementary constraints, and thus is feasible and optimal for (9). \blacksquare

Remark 3 *Observe that describing $\text{cl conv}(Q_T^{r1})$ requires two nonlinear inequalities in the original space of variables. More compactly, we can specify $\text{cl conv}(Q_T^{r1})$ using a single convex inequality, as*

$$\text{cl conv}(Q_T^{r1}) = \left\{ (\mathbf{z}, \boldsymbol{\beta}, t) \in [0, 1]^{|T|} \times \mathbb{R}^{|T|} \times \mathbb{R}_+ : \frac{(\mathbf{a}_T^\top \boldsymbol{\beta})^2}{\min\{1, z(T)\}} \leq t \right\}.$$

Finally, we point out that $\text{cl conv}(Q_T^{r1})$ is conic quadratic representable, as $(\mathbf{z}, \boldsymbol{\beta}, t) \in \text{cl conv}(Q_T^{r1})$ if and only if there exists w such that the system

$$\mathbf{z} \in [0, 1]^{|T|}, \boldsymbol{\beta} \in \mathbb{R}^{|T|}, t \in \mathbb{R}_+, w \in \mathbb{R}_+, w \leq 1, w \leq z(T), (\mathbf{a}_T^\top \boldsymbol{\beta})^2 \leq tw$$

is feasible, where the last constraint is a rotated conic quadratic constraint and all other constraints are linear.

2.2 General Case

Now consider again the mixed-integer optimization (3)

$$\mathbf{y}^\top \mathbf{y} + \min_{\boldsymbol{\beta}, \mathbf{z}, \mathbf{u}} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mu \left(\mathbf{e}^\top \mathbf{u} \right) + t \tag{13a}$$

$$\text{s.t. } \boldsymbol{\beta}^\top \left(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \right) \boldsymbol{\beta} \leq t \tag{13b}$$

$$\mathbf{e}^\top \mathbf{z} \leq k \tag{13c}$$

$$\boldsymbol{\beta} \leq \mathbf{u}, -\boldsymbol{\beta} \leq \mathbf{u} \tag{13d}$$

$$\boldsymbol{\beta} \circ (\mathbf{e} - \mathbf{z}) = \mathbf{0} \tag{13e}$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{z} \in \{0, 1\}^p, \mathbf{u} \in \mathbb{R}_+^p, t \in \mathbb{R} \tag{13f}$$

where the nonlinear terms of the objective is moved to constraint (13b). A direct application of (7) yields the inequality $\boldsymbol{\beta}^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \boldsymbol{\beta} \leq tz(P)$, which is weak and has no effect

when $z(P) \geq 1$. Instead, a more effective approach is to decompose the matrix $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ into a sum of low-dimensional rank-one matrices, and use inequality (7) to strengthen each quadratic term in the decomposition separately, as illustrated in Example 1 below.

Example 1 Consider the example with $p = 3$ and $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} = \begin{pmatrix} 25 & 15 & -5 \\ 15 & 18 & 0 \\ -5 & 0 & 11 \end{pmatrix}$. Then,

it follows that

$$\boldsymbol{\beta}^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \boldsymbol{\beta} = (5\beta_1 + 3\beta_2 - \beta_3)^2 + (3\beta_2 + \beta_3)^2 + 9\beta_3^2$$

and we have the corresponding valid inequality

$$\frac{(5\beta_1 + 3\beta_2 - \beta_3)^2}{\min\{1, z_1 + z_2 + z_3\}} + \frac{(3\beta_2 + \beta_3)^2}{\min\{1, z_2 + z_3\}} + 9\frac{\beta_3^2}{z_3} \leq t. \quad (14)$$

The decomposition of $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ illustrated in Example 1 is not unique. Since one does not obtain a strengthening when the denominator is one, it is important to have decomposition both rank-one and sparse. This motivates the question of how to find a decomposition that results in the best convex relaxation, i.e., that maximizes the left-hand side of (14). Specifically, let $\mathcal{P} \subseteq 2^P$ be a subset of the power set of P , i.e.,

$$\mathcal{P} = \{T_1, \dots, T_m\}$$

with $T_h \subseteq P$, $h = 1, \dots, m$. For each h , define a matrix variable \mathbf{A}_h whose nonzero elements correspond to the submatrix induced by T_h , and consider the valid inequality $\phi_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta}) \leq t$, where $\phi_{\mathcal{P}} : [0, 1]^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is defined as

$$\phi_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta}) := \max_{\mathbf{A}_h, \mathbf{R}} \boldsymbol{\beta}^\top \mathbf{R} \boldsymbol{\beta} + \sum_{h=1}^m \frac{\boldsymbol{\beta}^\top \mathbf{A}_h \boldsymbol{\beta}}{\min\{1, z(T_h)\}} \quad (15a)$$

$$\text{s.t. } \sum_{h=1}^m \mathbf{A}_h + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \quad (15b)$$

$$(\mathbf{A}_h)_{ij} = 0 \quad \forall h = 1, \dots, m, i \notin T_h \text{ or } j \notin T_h \quad (15c)$$

$$\mathbf{A}_h \in \mathcal{S}_+^P \quad \forall h = 1, \dots, m \quad (15d)$$

$$\mathbf{R} \in \mathcal{S}_+^P, \quad (15e)$$

where strengthening (7) is applied to each low-dimensional quadratic term $\boldsymbol{\beta}^\top \mathbf{A}_h \boldsymbol{\beta}$. For a fixed value of $(\mathbf{z}, \boldsymbol{\beta})$, problem (15) finds the best decomposition of the matrix $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ as a sum of positive semidefinite matrices \mathbf{A}_h , $h = 1, \dots, m$, and a remainder positive semidefinite matrix \mathbf{R} to maximize the strengthening.

For a given decomposition, the objective (15a) is convex in $(\mathbf{z}, \boldsymbol{\beta})$, thus $\phi_{\mathcal{P}}$ is a supremum of convex functions and is convex on its domain. Observe that the inclusion or omission of the empty set does not affect function $\phi_{\mathcal{P}}$, and we assume for simplicity that $\emptyset \in \mathcal{P}$.

Since inequalities (7) are ideal for rank-one matrices, inequality $\phi_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta}) \leq t$ is particularly strong if matrices \mathbf{A}_h are rank-one in optimal solutions of (15). As we now show, this is indeed the case if \mathcal{P} is downward closed.

Proposition 4 *If \mathcal{P} is downward closed, i.e., $V \in \mathcal{P} \implies U \in \mathcal{P}$ for all $U \subseteq V$, then there exists an optimal solution to (15) where all matrices \mathbf{A}_h are rank-one.*

Proof Let $T \in \mathcal{P}$, let \mathbf{A}_T be the matrix variable associated with T , and suppose \mathbf{A}_T is not rank-one in an optimal solution to (15), also suppose for simplicity that $T = \{1, \dots, p_0\}$ for some $p_0 \leq p$, and let $\bar{T}_i = \{i, \dots, p_0\}$ for $i = 1, \dots, p_0$. Since \mathbf{A}_T is positive semidefinite, there exists a Cholesky decomposition $\mathbf{A}_T = \mathbf{L}\mathbf{L}^\top$ where \mathbf{L} is a lower triangular matrix (possibly with zeros on the diagonal if \mathbf{A}_T is not positive definite). Let \mathbf{L}_i denote the i -th column of \mathbf{L} . Since \mathbf{A}_T is not a rank-one matrix, there exist at least two non-zero columns of \mathbf{L} . Let \mathbf{L}_j with $j > 1$ be the second non-zero column. Then

$$\begin{aligned} \frac{\beta_T^\top \mathbf{A}_T \beta_T}{\min\{1, z(T)\}} &= \frac{\beta_T^\top \left(\sum_{i \neq j} (\mathbf{L}_i \mathbf{L}_i^\top) \right) \beta_T}{\min\{1, z(T)\}} + \frac{\beta_T^\top (\mathbf{L}_j \mathbf{L}_j^\top) \beta_T}{\min\{1, z(T)\}} \\ &\leq \frac{\beta_T^\top \left(\sum_{i \neq j} (\mathbf{L}_i \mathbf{L}_i^\top) \right) \beta_T}{\min\{1, z(T)\}} + \frac{\beta_T^\top (\mathbf{L}_j \mathbf{L}_j^\top) \beta_T}{\min\{1, z(\bar{T}_j)\}}. \end{aligned} \quad (16)$$

Finally, since $\bar{T}_j \in \mathcal{P}$, the (better) decomposition (16) is feasible for (15), and the proposition is proven. \blacksquare

By dropping the complementary constraints (13e), replacing the integrality constraints $z \in \{0, 1\}^p$ with bound constraints $z \in [0, 1]^p$, and utilizing the convex function $\phi_{\mathcal{P}}$ to reformulate (13b), we obtain the convex relaxation of (1)

$$\mathbf{y}^\top \mathbf{y} + \min_{\beta, z, \mathbf{u}} -2\mathbf{y}^\top \mathbf{X}\beta + \mu \left(\mathbf{e}^\top \mathbf{u} \right) + \phi_{\mathcal{P}}(z, \beta) \quad (17a)$$

$$\mathbf{e}^\top z \leq k \quad (17b)$$

$$\beta \leq \mathbf{u}, -\beta \leq \mathbf{u} \quad (17c)$$

$$\beta \in \mathbb{R}^p, z \in [0, 1]^p, \mathbf{u} \in \mathbb{R}_+^p \quad (17d)$$

for a given $\mathcal{P} \subseteq 2^P$. In the next section, we give an interpretation of formulation (17) as a sparsity-inducing regularization penalty.

2.3 Interpretation as Regularization

Note that the relaxation (17) can be rewritten as:

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_2^2 + \mu \|\beta\|_1 + \rho_{R1}(\beta; k)$$

where

$$\rho_{R1}(\beta; k) := \min_{z \in [0, 1]^p} \phi_{\mathcal{P}}(z, \beta) - \beta^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \beta \text{ s.t. } \mathbf{e}^\top z \leq k. \quad (18)$$

is the (non-convex) *rank-one regularization penalty*. Observe that $\rho_{R1}(\beta; k)$ is the difference of two convex functions: the quadratic function $\beta^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \beta$ arising from the fitness term and the Tikhonov regularization; and the projection of its convexification $\phi_{\mathcal{P}}(z, \beta)$ in the original space of the regression variables β . As we now show, unlike the usual ℓ_1 penalty, the *rank-one regularization penalty* does not induce a bias when β is sparse.

Theorem 5 *If $\|\beta\|_0 \leq k$, then $\rho_{R1}(\beta; k) = 0$.*

Proof Let $(\beta, z) \in \mathbb{R}^p \times [0, 1]^p$, and let \mathbf{R} and \mathbf{A}_h , $h = 1, \dots, m$, correspond to an optimal solution of (15). Since

$$\begin{aligned} \beta^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \beta &= \beta^\top \mathbf{R} \beta + \sum_{h=1}^m \beta^\top \mathbf{A}_h \beta \\ &\leq \beta^\top \mathbf{R} \beta + \sum_{h=1}^m \frac{\beta^\top \mathbf{A}_h \beta}{\min\{1, z(T_h)\}} = \phi_{\mathcal{P}}(z, \beta), \end{aligned}$$

it follows that $\rho_{R1}(\beta; k) \geq 0$ for any $\beta \in \mathbb{R}^p$. Now let $\hat{\beta}$ satisfy $\|\hat{\beta}\|_0 \leq k$, let $\hat{T} = \{i \in P : \hat{\beta}_i \neq 0\}$ be the support of $\hat{\beta}$ and let \hat{z} such that $\hat{z}_i = \mathbb{1}_{i \in \hat{T}}$ be the indicator vector of \hat{T} . By construction, $e^\top \hat{z} \leq k$ and \hat{z} is feasible for problem (18). Moreover

$$\begin{aligned} \rho_{R1}(\hat{\beta}; k) &\leq \phi_{\mathcal{P}}(\hat{z}, \hat{\beta}) - \hat{\beta}^\top (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}) \hat{\beta} \\ &= \sum_{\substack{1 \leq h \leq m \\ T_h \cap \hat{T} \neq \emptyset}} \left(\frac{\hat{\beta}^\top \mathbf{A}_h \hat{\beta}}{\min\{1, \hat{z}(T_h)\}} - \hat{\beta}^\top \mathbf{A}_h \hat{\beta} \right) = 0. \end{aligned}$$

Thus, $\rho_{R1}(\hat{\beta}; k) = 0$. ■

The rank-one regularization penalty ρ_{R1} can also be interpreted from an optimization perspective: note that problem (15) is the *separation* problem that, given $(\beta, z) \in \mathbb{R}^p \times [0, 1]^p$, finds a decomposition that results in a most violated inequality after applying the rank-one strengthening. Thus, the regularization penalty $\rho_{R1}(\beta; k)$ is precisely the violation of this inequality when z is chosen optimally.

In §3 we derive an explicit form of $\rho_{R1}(\beta; k)$ when $p = 2$; Figure 1 plots the graphs of the usual regularization penalties and ρ_{R1} for the two-dimensional case, and Figure 2 illustrates the better sparsity inducing properties of regularization ρ_{R1} . Deriving explicit forms of ρ_{R1} is cumbersome for $p \geq 3$. Fortunately, problem (17) can be explicitly reformulated in an extended space as an SDP and tackled using off-the-shelf conic optimization solvers.

2.4 Extended SDP Formulation

To state the extended SDP formulation, in addition to variables $z \in [0, 1]^p$ and $\beta \in \mathbb{R}^p$, we introduce variables $w \in [0, 1]^m$ corresponding to terms $w_h := \min\{1, z(T_h)\}$ and $\mathbf{B} \in \mathbb{R}^{p \times p}$ corresponding to terms $B_{ij} = \beta_i \beta_j$.

Theorem 6 *Problem (17) is equivalent to the SDP*

$$\mathbf{y}^\top \mathbf{y} + \min -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (19a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{z} \leq k \quad (19b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, \quad -\boldsymbol{\beta} \leq \mathbf{u} \quad (19c)$$

$$w_h \leq \mathbf{e}_{T_h}^\top \mathbf{z}_{T_h} \quad \forall h = 1, \dots, m \quad (19d)$$

$$w_h \mathbf{B}_{T_h} - \boldsymbol{\beta}_{T_h} \boldsymbol{\beta}_{T_h}^\top \in \mathcal{S}_+^{T_h} \quad \forall h = 1, \dots, m \quad (19e)$$

$$\mathbf{B} - \boldsymbol{\beta} \boldsymbol{\beta}^\top \in \mathcal{S}_+^P \quad (19f)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \quad \mathbf{z} \in [0, 1]^p, \quad \mathbf{u} \in \mathbb{R}_+^p, \quad \mathbf{w} \in [0, 1]^m, \quad \mathbf{B} \in \mathbb{R}^{p \times p}. \quad (19g)$$

Observe that (19) is indeed an SDP, as

$$w_h \mathbf{B}_{T_h} - \boldsymbol{\beta}_{T_h} \boldsymbol{\beta}_{T_h}^\top \in \mathcal{S}_+^{T_h} \Leftrightarrow \begin{pmatrix} w_h & \boldsymbol{\beta}_{T_h}^\top \\ \boldsymbol{\beta}_{T_h} & \mathbf{B}_{T_h} \end{pmatrix} \succeq 0. \quad (20)$$

Indeed, from Schur's complement, the right-hand side of (20) is equivalent to $w_h \geq 0$ (automatically satisfied) and $\mathbf{B}_{T_h} - \frac{1}{w_h} \boldsymbol{\beta}_{T_h} \boldsymbol{\beta}_{T_h}^\top \succeq 0$. Similarly, constraint (19f) can be modeled as $\begin{pmatrix} 1 & \boldsymbol{\beta}^\top \\ \boldsymbol{\beta} & \mathbf{B} \end{pmatrix} \succeq 0$. Thus, constraints (19e) and (19f) are indeed SDP-representable, and the remaining constraints and objective are linear.

Proof It is easy to check that (19) is strictly feasible (set $\boldsymbol{\beta} = 0$, $\mathbf{z} = \mathbf{e}$, $\mathbf{w} > \mathbf{0}$ and $\mathbf{B} = \mathbf{I}$). Adding surplus variables $\boldsymbol{\Gamma}$, $\boldsymbol{\Gamma}_h$ for $h = 1, \dots, m$, write (19) as

$$\begin{aligned} \mathbf{y}^\top \mathbf{y} + \min_{(\boldsymbol{\beta}, \mathbf{z}, \mathbf{u}, \mathbf{w}) \in C} \left\{ -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \min_{\mathbf{B}, \boldsymbol{\Gamma}_h, \boldsymbol{\Gamma}} \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \right\} \\ \text{s.t. } w_h \mathbf{B}_{T_h} - \boldsymbol{\Gamma}_h = \boldsymbol{\beta}_{T_h} \boldsymbol{\beta}_{T_h}^\top \quad \forall h \quad (\mathbf{A}_h) \\ \mathbf{B} - \boldsymbol{\Gamma} = \boldsymbol{\beta} \boldsymbol{\beta}^\top \quad (\mathbf{R}) \\ \boldsymbol{\Gamma}_h \in \mathcal{S}_+^{T_h} \quad \forall h \\ \boldsymbol{\Gamma} \in \mathcal{S}_+^P \\ \mathbf{B} \in \mathbb{R}^{p \times p}, \end{aligned}$$

where $C = \{\boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{z} \in [0, 1]^p, \mathbf{u} \in \mathbb{R}_+^p, \mathbf{w} \in [0, 1]^m : (19b), (19c), (19d)\}$. Using conic duality for the inner minimization problem, we find the dual

$$\begin{aligned} \mathbf{y}^\top \mathbf{y} + \min_{(\boldsymbol{\beta}, \mathbf{z}, \mathbf{u}, \mathbf{w}) \in C} \left\{ -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \max_{\mathbf{A}_h, \mathbf{R}} \langle \boldsymbol{\beta} \boldsymbol{\beta}^\top, \mathbf{R} \rangle + \sum_{h=1}^m \langle \boldsymbol{\beta} \boldsymbol{\beta}^\top, \mathbf{A}_h \rangle \right\} \\ \text{s.t. } \sum_{j=1}^m w_j \mathbf{A}_j + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \\ (\mathbf{A}_h)_{ij} = 0 \quad \text{for } i \notin T_h \text{ or } j \notin T_h \\ \mathbf{A}_h \in \mathcal{S}_+^P \quad \forall h \in \mathcal{P} \\ \mathbf{R} \in \mathcal{S}_+^P. \end{aligned}$$

After substituting $\bar{\mathbf{A}}_h = w_h \mathbf{A}_h$ and noting that there exists an optimal solution with $w_h = \min\{1, z(T_h)\}$, we obtain formulation (15). \blacksquare

Note that if $\mathcal{P} = \{\emptyset\}$, there is no strengthening and (19) is equivalent to **elastic net** ($\lambda, \mu > 0$), **lasso** ($\lambda = 0, \mu > 0$), **ridge regression** ($\lambda > 0, \mu = 0$) or **ordinary least squares** ($\lambda = \mu = 0$). As $|\mathcal{P}|$ increases, the quality of the conic relaxation (19) for the non-convex ℓ_0 -problem (1) improves, but the computational burden required to solve the resulting SDP also increases. In particular, the *full* rank-one strengthening with $\mathcal{P} = 2^P$ requires 2^P semidefinite constraints and is impractical. Proposition 4 suggests using down-monotone sets \mathcal{P} with limited size

$$\mathbf{y}^\top \mathbf{y} + \min -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (21a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{z} \leq k \quad (21b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, -\boldsymbol{\beta} \leq \mathbf{u} \quad (21c)$$

$$(\text{sdp}_r) \quad 0 \leq w_T \leq \min\{1, \mathbf{e}_T^\top \mathbf{z}_T\} \quad \forall T : |T| \leq r \quad (21d)$$

$$w_T \mathbf{B}_T - \boldsymbol{\beta}_T \boldsymbol{\beta}_T^\top \in \mathcal{S}_+^T \quad \forall T : |T| \leq r \quad (21e)$$

$$\mathbf{B} - \boldsymbol{\beta} \boldsymbol{\beta}^\top \in \mathcal{S}_+^P \quad (21f)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{z} \in [0, 1]^p, \mathbf{u} \in \mathbb{R}_+^p, \mathbf{B} \in \mathbb{R}^{p \times p}, \mathbf{w} \in \mathbb{R}^m \quad (21g)$$

for some $r \in \mathbb{Z}_+$, where $m = |\{T \subseteq P : |T| \leq r\}|$. Note that in the above formulation, w_T is a scalar corresponding to the T -th coordinate of the m -dimensional vector \mathbf{w} . If $r = 1$, then sdp_1 reduces to the formulation of the optimal **perspective relaxation** proposed by Dong et al. (2015), which is equivalent to using MC_+ regularization. Our computations experiments show that whereas sdp_1 may be a weak convex relaxation for problems with low diagonal dominance, sdp_2 achieves excellent relaxation bounds even for the case of low diagonal dominance within reasonable computing times. For clarity, we give the explicit form of the case sdp_2 :

$$\mathbf{y}^\top \mathbf{y} + \min -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (22a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{z} \leq k \quad (22b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, -\boldsymbol{\beta} \leq \mathbf{u} \quad (22c)$$

$$\begin{pmatrix} z_i & \beta_i \\ \beta_i & B_{ii} \end{pmatrix} \succeq 0 \quad \forall i = 1, \dots, p \quad (22d)$$

$$(\text{sdp}_2) \quad 0 \leq w_{ij} \leq \min\{1, z_i + z_j\} \quad \forall i < j \quad (22e)$$

$$\begin{pmatrix} w_{ij} & \beta_i & \beta_j \\ \beta_i & B_{ii} & B_{ij} \\ \beta_j & B_{ij} & B_{jj} \end{pmatrix} \succeq 0 \quad \forall i < j \quad (22f)$$

$$\begin{pmatrix} 1 & \boldsymbol{\beta}^\top \\ \boldsymbol{\beta} & \mathbf{B} \end{pmatrix} \succeq 0 \quad (22g)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{z} \in [0, 1]^p, \mathbf{u} \in \mathbb{R}_+^p, \mathbf{B} \in \mathbb{R}^{p \times p}, \mathbf{w} \in \mathbb{R}^{p(p-1)/2}. \quad (22h)$$

Constraints (22e) and (22f) correspond exactly to constraints (21d)-(21e) for the cases with $|T| = 2$. Moreover, for cases where T is a singleton, that is, $T = \{i\}$ for some $i \in \{1, \dots, p\}$, constraints $0 \leq w_i \leq z_i$ can be omitted since it can be easily verified that

$w_i = z_i$ in any optimal solution; thus constraints (21e) reduce, after substitution of w_i , to constraints (22d). Finally, constraints (22b), (22c) and (22g) correspond to constraints (21b), (21c) and (21f), respectively.

2.5 Additional Comments

Just like `lasso` can be interpreted as the best possible convex relaxation obtained from the set with bounded continuous variables $\{z \in \{0, 1\}, \beta \in [0, 1] : \beta(1 - z) = 0\}$ and the `perspective relaxation` is the best possible relaxation obtained from set Q_1 , the relaxations proposed in this section are the best possible convex relaxations obtained from study of Q_T^{r1} . Since Q_T^{r1} generalizes these two simpler sets, it follows that the proposed formulations are a better approximation of (3) than `lasso` and the `perspective relaxation`. Nonetheless, Q_T^{r1} is still considerably simpler than the feasible region of (3) – which involves constraints on the binary variables and general convex quadratic functions as described in set Q_T . We now briefly discuss two recent results that shed additional light on the strength of the formulations.

Wei et al. (2022) consider a generalization of Q_T^{r1} in which the binary variables are subject to additional constraints. The authors find that with the k -sparsity constraint $\sum_{i=1}^p z_i \leq k$, the rank-one relaxation described in Theorem 2 is still the best possible formulation. Moreover, the authors show that the convex hull for other classes of constraints has a similar structure to $\text{cl conv}(Q_T^{r1})$, and formulation sdp_r can be extended to those cases.

Wei et al. (2024) study set Q_T , and show that $\text{cl conv}(Q_T)$ can be described in an extended formulation with $O(p^2)$ additional variables as the intersection of a single conic constraint and a polyhedron Ψ . The proof of this result is not constructive: an explicit description of Ψ is not given, as it requires an exponential number of linear inequalities. Nonetheless, the authors show that formulations sdp_r can be interpreted as relaxations in this extended space obtained by adding linear inequalities that are *guaranteed* to define high-dimensional faces of Ψ – the dimension of the face is larger for small values of r , providing additional theoretical justification that the nonlinear inequalities (22d) and (22f) are in fact strong approximations of $\text{cl conv}(Q_T)$.

3 Regularization for the Two-Dimensional Case

To better understand the properties of the proposed conic relaxations, in this section, we study them from a regularization perspective. Consider formulation (17b) in Lagrangean form with multiplier κ :

$$\mathbf{y}^\top \mathbf{y} + \min - 2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \phi_p(\mathbf{z}, \boldsymbol{\beta}) + \kappa \mathbf{e}^\top \mathbf{z} \quad (23a)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, \quad -\boldsymbol{\beta} \leq \mathbf{u} \quad (23b)$$

$$\boldsymbol{\beta} \in \mathbb{R}^P, \quad \mathbf{z} \in [0, 1]^P, \quad \mathbf{u} \in \mathbb{R}_+^P, \quad (23c)$$

where $p = 2$, and

$$\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} = \begin{pmatrix} 1 + \delta_1 & 1 \\ 1 & 1 + \delta_2 \end{pmatrix}. \quad (24)$$

Observe that assumption (24) is without loss of generality, provided that $\mathbf{X}^\top \mathbf{X}$ is not diagonal: given a two-dimensional convex quadratic function $a_1\beta_1^2 + 2a_{12}\beta_1\beta_2 + a_2\beta_2^2$ (with

$a_{12} \neq 0$), the substitution $\bar{\beta}_1 = \alpha\beta_1$ and $\bar{\beta}_2 = (a_{12}/\alpha)\beta_2$ with $|a_{12}|/a_2 \leq \alpha \leq a_1$ yields a quadratic form satisfying (24). Also note that we are using the Lagrangian form instead of the cardinality constrained form given in (18) for simplicity; however, since $\phi_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta})$ is convex in \mathbf{z} , there exists a value of κ such that both forms are equivalent, i.e., result in the same optimal solutions $\hat{\boldsymbol{\beta}}$ for the regression problem, and the objective values differ by the constant $\kappa \cdot k$.

If $\mathcal{P} = \{\emptyset, \{1\}, \{2\}\}$, then (23) reduces to a perspective strengthening of the form

$$\mathbf{y}'\mathbf{y} + \min_{\mathbf{z} \in [0,1]^2, \boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}'\mathbf{X}\boldsymbol{\beta} + (\beta_1 + \beta_2)^2 + \delta_1 \frac{\beta_1^2}{z_1} + \delta_2 \frac{\beta_2^2}{z_2} + \mu\|\boldsymbol{\beta}\|_1 + \kappa\|\mathbf{z}\|_1. \quad (25)$$

The links between (25) and regularization were studied³ in Dong et al. (2015).

Proposition 7 (Dong et al. (2015)) *Problem (25) is equivalent to the regularization problem*

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2 + \mu\|\boldsymbol{\beta}\|_1 + \rho_{MC_+}(\boldsymbol{\beta}; \kappa, \boldsymbol{\delta})$$

where

$$\rho_{MC_+}(\boldsymbol{\beta}; \kappa, \boldsymbol{\delta}) = \begin{cases} \sum_{i=1}^2 (2\sqrt{\kappa\delta_i}|\beta_i| - \delta_i\beta_i^2) & \text{if } \delta_i\beta_i^2 \leq \kappa, i = 1, 2 \\ \kappa + 2\sqrt{\kappa\delta_i}|\beta_i| - \delta_i\beta_i^2 & \text{if } \delta_i\beta_i^2 \leq \kappa \text{ and } \delta_j\beta_j^2 > \kappa \\ 2\kappa & \text{if } \delta_i\beta_i^2 > \kappa, i = 1, 2. \end{cases}$$

Regularization ρ_{MC_+} is non-convex and separable. Moreover, as pointed out by Dong et al. (2015), the regularization given in Proposition 7 is the same as the Minimax Concave Penalty given by Zhang et al. (2010); and, if $\lambda = \delta_1 = \delta_2$, then the regularization given in Proposition 7 reduces to the reverse Huber penalty derived by Pilanci et al. (2015). Observe that the regularization function ρ_{MC_+} is highly dependent on the diagonal dominance $\boldsymbol{\delta}$: specifically, in the low diagonal dominance setting with $\boldsymbol{\delta} = \mathbf{0}$, we find that $\rho_{MC_+}(\boldsymbol{\beta}; \kappa, \mathbf{0}) = 0$.

We now consider conic formulation (23) for the case $\mathcal{P} = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$, corresponding to the full rank-one strengthening:

$$\mathbf{y}^\top \mathbf{y} + \min_{\mathbf{z} \in [0,1]^2, \boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \frac{(\beta_1 + \beta_2)^2}{\min\{1, z_1 + z_2\}} + \delta_1 \frac{\beta_1^2}{z_1} + \delta_2 \frac{\beta_2^2}{z_2} + \mu\|\boldsymbol{\beta}\|_1 + \kappa\|\mathbf{z}\|_1. \quad (26)$$

Proposition 8 *Problem (26) is equivalent to the regularization problem*

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2 + \mu\|\boldsymbol{\beta}\|_1 + \rho_{R1}(\boldsymbol{\beta}; \kappa, \boldsymbol{\delta})$$

3. The case with $\mu = 0$ is explicitly considered in Dong et al. (2015), but the results extend straightforwardly to the case with $\mu > 0$. The results presented here differ slightly from those by Dong et al. (2015) to account for a different scaling in the objective function.

where

$$\rho_{RI}(\boldsymbol{\beta}; \kappa, \boldsymbol{\delta}) = \begin{cases} 2\sqrt{\kappa}\sqrt{\boldsymbol{\beta}'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\boldsymbol{\beta} + 2\sqrt{\delta_1\delta_2}|\beta_1\beta_2|} - \boldsymbol{\beta}'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\boldsymbol{\beta} & \text{if } \boldsymbol{\beta}'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\boldsymbol{\beta} + 2\sqrt{\delta_1\delta_2}|\beta_1\beta_2| < \kappa \\ \kappa + 2\sqrt{\delta_1\delta_2}|\beta_1\beta_2| & \\ \sum_{i=1}^2 (2\sqrt{\kappa\delta_i}|\beta_i| - \delta_i\beta_i^2) & \text{if } (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2 \leq \kappa \leq \boldsymbol{\beta}'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\boldsymbol{\beta} + 2\sqrt{\delta_1\delta_2}|\beta_1\beta_2| \\ \kappa + \sqrt{\kappa\delta_i}|\beta_i| - \delta_i\beta_i^2 & \text{if } (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2 > \kappa \text{ \& } \delta_i\beta_i^2 \leq \kappa, i = 1, 2 \\ 2\kappa & \text{if } \delta_i\beta_i^2 > \kappa, i = 1, 2. \end{cases}$$

Proof We prove the result by projecting out the z variables in (26), i.e., giving closed-form solutions for them. There are three cases to consider, depending on the optimal value for $z_1 + z_2$.

• **Case 1:** $z_1 + z_2 < 1$ In this case, we find by setting the derivatives of the objective in (26) with respect to z_1 and z_2 that

$$\left. \begin{aligned} \kappa - \delta_1 \frac{\beta_1^2}{z_1^2} - \frac{(\beta_1 + \beta_2)^2}{(z_1 + z_2)^2} &= 0 \\ \kappa - \delta_2 \frac{\beta_2^2}{z_2^2} - \frac{(\beta_1 + \beta_2)^2}{(z_1 + z_2)^2} &= 0 \end{aligned} \right\} \implies z_2 = \sqrt{\frac{\delta_2}{\delta_1}} \frac{|\beta_2|}{|\beta_1|} z_1.$$

Define $\bar{z} := \frac{z_1}{\sqrt{\delta_1}|\beta_1|}$, so $z_2 = \sqrt{\delta_2}|\beta_2|\bar{z}$, and $z_1 + z_2 = (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)\bar{z}$. Moreover, we find that (26) reduces to

$$\begin{aligned} \mathbf{y}^\top \mathbf{y} + \min_{\bar{z} > 0, \boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mu\|\boldsymbol{\beta}\|_1 \\ + \frac{(\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2}{(\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)\bar{z}} + \kappa (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)\bar{z}. \end{aligned} \quad (27)$$

An optimal solution of (27) is attained at

$$\bar{z}^* = \sqrt{\frac{(\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2}{\kappa (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)}} = \frac{\sqrt{(\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2}}{\sqrt{\kappa} (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)}$$

with objective value

$$\begin{aligned} &\mathbf{y}^\top \mathbf{y} + \min_{\boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mu\|\boldsymbol{\beta}\|_1 + 2\sqrt{\kappa}\sqrt{(\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2} \\ &= \min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2 + \mu\|\boldsymbol{\beta}\|_1 \\ &\quad + \left(2\sqrt{\kappa}\sqrt{(\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2} - (\beta_1 + \beta_2)^2 - \delta_1\beta_1^2 - \delta_2\beta_2^2 \right). \end{aligned}$$

Finally, this case happens when $z_1 + z_2 < 1 \Leftrightarrow (\beta_1 + \beta_2)^2 + (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2 < \kappa$.

• **Case 2:** $z_1 + z_2 > 1$ In this case, we find by setting the derivatives of the objective in (26) with respect to z_1 and z_2 that $\bar{z}_i = \sqrt{\frac{\delta_i}{\kappa}}|\beta_i|$ for $i = 1, 2$. Thus, in this case, for an optimal solution \mathbf{z}^* of (26), we have $z_i^* = \min\{\bar{z}_i, 1\}$, and problem (26) reduces to

$$\begin{aligned} & \mathbf{y}^\top \mathbf{y} + \min_{\boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + (\beta_1 + \beta_2)^2 + \sum_{i=1}^2 \max\left\{\delta_i \beta_i^2, \sqrt{\kappa \delta_i} |\beta_i|\right\} + \mu \|\boldsymbol{\beta}\|_1 \\ & \quad + \sum_{i=1}^2 \min\left\{\sqrt{\kappa \delta_i} |\beta_i|, \kappa\right\} \\ & = \min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 + \mu \|\boldsymbol{\beta}\|_1 + \sum_{i=1}^2 \left(\max\left\{\delta_i \beta_i^2, \sqrt{\kappa \delta_i} |\beta_i|\right\} + \min\left\{\sqrt{\kappa \delta_i} |\beta_i|, \kappa\right\} - \delta_i \beta_i^2\right) \\ & = \min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 + \mu \|\boldsymbol{\beta}\|_1 + \begin{cases} \sum_{i=1}^2 (2\sqrt{\kappa \delta_i} |\beta_i| - \delta_i \beta_i^2) & \text{if } \delta_i \beta_i^2 \leq \kappa, i = 1, 2 \\ \sqrt{\kappa \delta_i} |\beta_i| - \delta_i \beta_i^2 + \kappa & \text{if } \delta_i \beta_i^2 \leq \kappa \text{ \& } \delta_j \beta_j^2 > \kappa \\ 2\kappa & \text{if } \delta_i \beta_i^2 > \kappa, i = 1, 2. \end{cases} \end{aligned}$$

Finally, this case happens when $z_1 + z_2 > 1 \Leftrightarrow (\sqrt{\delta_1}|\beta_1| + \sqrt{\delta_2}|\beta_2|)^2 > \kappa$. Observe that, in this case, the penalty function is precisely the one given in Proposition 7.

• **Case 3:** $z_1 + z_2 = 1$ In this case, problem (26) reduces to

$$\mathbf{y}^\top \mathbf{y} + \min_{0 \leq z_1 \leq 1, \boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + (\beta_1 + \beta_2)^2 + \delta_1 \frac{\beta_1^2}{z_1} + \delta_2 \frac{\beta_2^2}{1 - z_1} + \mu \|\boldsymbol{\beta}\|_1 + \kappa. \quad (28)$$

Setting derivative with respect to z_1 in (28) to 0, we have

$$\begin{aligned} 0 & = \delta_1 \beta_1^2 (1 - z_1)^2 - \delta_2 \beta_2^2 z_1^2 \\ & = \delta_1 \beta_1^2 - 2\delta_1 \beta_1^2 z_1 + (\delta_1 \beta_1^2 - \delta_2 \beta_2^2) z_1^2. \end{aligned}$$

Thus, we find that

$$\begin{aligned} z_1 & = \frac{2\delta_1 \beta_1^2 \pm \sqrt{4\delta_1^2 \beta_1^4 - 4\delta_1 \beta_1^2 (\delta_1 \beta_1^2 - \delta_2 \beta_2^2)}}{2(\delta_1 \beta_1^2 - \delta_2 \beta_2^2)} \\ & = \frac{\delta_1 \beta_1^2 \pm \sqrt{\delta_1 \delta_2} |\beta_1 \beta_2|}{\delta_1 \beta_1^2 - \delta_2 \beta_2^2} = \frac{\sqrt{\delta_1} |\beta_1| (\sqrt{\delta_1} |\beta_1| \pm \sqrt{\delta_2} |\beta_2|)}{(\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|)(\sqrt{\delta_1} |\beta_1| - \sqrt{\delta_2} |\beta_2|)}. \end{aligned}$$

Moreover, since $0 \leq z_1 \leq 1$, we have $z_1 = \frac{\sqrt{\delta_1} |\beta_1|}{\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|}$ and $1 - z_1 = \frac{\sqrt{\delta_2} |\beta_2|}{\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|}$. Substituting in (28), we find the equivalent form

$$\begin{aligned} & \mathbf{y}^\top \mathbf{y} + \min_{\boldsymbol{\beta} \in \mathbb{R}^2} -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + (\beta_1 + \beta_2)^2 + \left(\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|\right)^2 + \mu \|\boldsymbol{\beta}\|_1 + \kappa \\ & = \min_{\boldsymbol{\beta} \in \mathbb{R}^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 + \mu \|\boldsymbol{\beta}\|_1 + \kappa + 2\sqrt{\delta_1 \delta_2} |\beta_1 \beta_2|. \end{aligned}$$

This final case occurs when neither case 1 or 2 does, i.e., when $(\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|)^2 \leq \kappa \leq (\beta_1 + \beta_2)^2 + (\sqrt{\delta_1} |\beta_1| + \sqrt{\delta_2} |\beta_2|)^2$. \blacksquare

Observe that, unlike ρ_{MC_+} , the function ρ_{R1} is not separable in β_1 and β_2 and does not vanish when $\delta = 0$: indeed, for $\delta = 0$ we find that

$$\rho_{\text{R1}}(\beta; \kappa, \mathbf{0}) = \begin{cases} 2\sqrt{\kappa}\sqrt{\beta'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\beta} - \beta'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\beta & \text{if } \beta'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\beta < \kappa \\ \kappa & \text{if } 0 \leq \kappa \leq \beta'(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})\beta. \end{cases}$$

The plots of ρ_{MC_+} and ρ_{R1} shown in Figures 1 and 2 correspond to setting the natural value $\kappa = 1$.

4 Conic Quadratic Relaxations

As mentioned in §1, strong convex relaxations of problem (1), such as sdp_τ , can either be directly used to obtain good estimators via conic optimization, which is the approach we use in our computations or can be embedded in a branch-and-bound algorithm to solve (1) to optimality. However, using SDP formulations such as (19) in branch-and-bound may be daunting since, to date, efficient branch-and-bound algorithms with SDP relaxations are not available. In contrast, conic quadratic optimization problems are considerably easier to solve than semidefinite optimization problems, thus scaling to larger dimensions. Moreover, there exist off-the-shelf mixed-integer conic quadratic optimization solvers that are actively maintained and improved by numerous software vendors. In this section, we show how the proposed conic relaxations, and specifically sdp_2 , can be implemented in a conic quadratic framework. The resulting convex formulations can then be directly used as a fast approximation to the SDP formulations presented in §2, and pave the way towards an integration with branch-and-bound solvers⁴.

4.1 Two-Dimensional PSD Constraints

Constraint (22d), $\beta_i^2 \leq z_i B_{ii}$, is a rotated cone constraint as $z_i \geq 0$ and $B_{ii} \geq 0$ in any feasible solution of (21), and thus conic quadratic representable.

4.2 Three-Dimensional PSD Constraints

As we now show, constraints (22f) can be accurately approximated using conic quadratic constraints.

4. An effective implementation would require careful constraint management strategies and integration with the different aspects of branch-and-bound solvers, e.g., branching strategies and heuristics. Such an implementation is beyond the scope of the paper.

Proposition 9 *Problem sdp_2 is equivalent to the optimization problem*

$$\mathbf{y}^\top \mathbf{y} + \min -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (29a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{z} \leq k \quad (29b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, \quad -\boldsymbol{\beta} \leq \mathbf{u} \quad (29c)$$

$$z_i B_{ii} \geq \beta_i^2 \quad \forall i \in P \quad (29d)$$

$$0 \leq w_{ij} \leq 1, \quad w_{ij} \leq z_i + z_j \quad \forall i \neq j \quad (29e)$$

$$0 \geq \max_{\alpha \geq 0} \left\{ \frac{\alpha \beta_i^2 + 2\beta_i \beta_j + \beta_j^2 / \alpha}{w_{ij}} - 2B_{ij} - \alpha B_{ii} - B_{jj} / \alpha \right\} \quad \forall i \neq j \quad (29f)$$

$$0 \geq \max_{\alpha \geq 0} \left\{ \frac{\alpha \beta_i^2 - 2\beta_i \beta_j + \beta_j^2 / \alpha}{w_{ij}} + 2B_{ij} - \alpha B_{ii} - B_{jj} / \alpha \right\} \quad \forall i \neq j \quad (29g)$$

$$\mathbf{B} - \boldsymbol{\beta}\boldsymbol{\beta}' \in \mathcal{S}_+^P \quad (29h)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \quad \mathbf{z} \in [0, 1]^p, \quad \mathbf{u} \in \mathbb{R}_+^p, \quad \mathbf{B} \in \mathbb{R}^{p \times p}. \quad (29i)$$

Proof It suffices to compute the optimal value of α in (29f)–(29g). Observe that the rhs of (29f) can be written as

$$v = \frac{2\beta_i \beta_j}{w_{ij}} - 2B_{ij} - \min_{\alpha \geq 0} \left\{ \alpha \left(B_{ii} - \frac{\beta_i^2}{w_{ij}} \right) + \frac{1}{\alpha} \left(B_{jj} - \frac{\beta_j^2}{w_{ij}} \right) \right\}. \quad (30)$$

Moreover, in an optimal solution of (29), we have that $w_{ij} = \min\{1, z_i + z_j\}$. Thus, due to constraints (29d), we find that $B_{ii} - \beta_i^2/w_{ij} \geq 0$ in optimal solutions of (29), and equality only occurs if either $z_i = 1$ or $z_j = 0$. If either $B_{ii} = \beta_i^2/\min\{1, z_i + z_j\}$ or $B_{jj} = \beta_j^2/\min\{1, z_i + z_j\}$, then the optimal value of (30) is $v = 2\beta_i \beta_j / \min\{1, z_i + z_j\} - 2B_{ij}$, by setting $\alpha \rightarrow \infty$ or $\alpha = 0$, respectively. Otherwise, the optimal α equals

$$\alpha = \sqrt{\frac{B_{jj} w_{ij} - \beta_j^2}{B_{ii} w_{ij} - \beta_i^2}}, \quad (31)$$

with the objective value

$$v = \frac{2\beta_i \beta_j}{w_{ij}} - 2B_{ij} - 2\sqrt{\left(B_{ii} - \frac{\beta_i^2}{w_{ij}} \right) \left(B_{jj} - \frac{\beta_j^2}{w_{ij}} \right)}.$$

Observe that this expression is also correct when $B_{ii} = \beta_i^2/\min\{1, z_i + z_j\}$ or $B_{jj} = \beta_j^2/\min\{1, z_i + z_j\}$. Thus, constraint (29f) reduces to

$$0 \geq \beta_i \beta_j - B_{ij} w_{ij} - \sqrt{(B_{ii} w_{ij} - \beta_i^2) (B_{jj} w_{ij} - \beta_j^2)}. \quad (32)$$

Similarly, it can be shown that constraint (29g) reduces to

$$0 \geq -\beta_i \beta_j + B_{ij} w_{ij} - \sqrt{(B_{ii} w_{ij} - \beta_i^2) (B_{jj} w_{ij} - \beta_j^2)}. \quad (33)$$

More compactly, constraints (32)–(33) are equivalent to

$$(w_{ij}B_{ii} - \beta_i^2)(w_{ij}B_{jj} - \beta_j^2) \geq (w_{ij}B_{ij} - \beta_i\beta_j)^2. \quad (34)$$

Moreover, note that constraints (21e) with $T = \{i, j\}$ are equivalent to

$$\begin{aligned} & \begin{pmatrix} w_{ij}B_{ii} - \beta_i^2 & w_{ij}B_{ij} - \beta_i\beta_j \\ w_{ij}B_{ij} - \beta_i\beta_j & w_{ij}B_{jj} - \beta_j^2 \end{pmatrix} \in \mathcal{S}_+^2 \\ & \Leftrightarrow w_{ij}B_{ii} - \beta_i^2 \geq 0, w_{ij}B_{jj} - \beta_j^2 \geq 0, \text{ and (34)}. \end{aligned}$$

Since the first two constraints are implied by (29d) and $w_{ij} = \min\{1, z_i + z_j\}$ in optimal solutions, the proof is complete. \blacksquare

Observe that, for any fixed value of α , constraints (29f)–(29g) are conic quadratic representable. Thus, we can obtain relaxations of (29) of the form

$$\mathbf{y}^\top \mathbf{y} + \min - 2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (35a)$$

$$\text{s.t. (29b), (29c), (29d), (29e), (29h), (29i)} \quad (35b)$$

$$0 \geq \frac{\alpha\beta_i^2 + 2\beta_i\beta_j + \beta_j^2/\alpha}{\min\{1, z_i + z_j\}} - 2B_{ij} - \alpha B_{ii} - B_{jj}/\alpha, \forall i \neq j, \alpha \in V_{ij}^+ \quad (35c)$$

$$0 \geq \frac{\alpha\beta_i^2 - 2\beta_i\beta_j + \beta_j^2/\alpha}{\min\{1, z_i + z_j\}} + 2B_{ij} - \alpha B_{ii} - B_{jj}/\alpha, \forall i \neq j, \alpha \in V_{ij}^-, \quad (35d)$$

where V_{ij}^+ and V_{ij}^- are any finite subsets of \mathbb{R}_+ . Relaxation (35) can be refined dynamically: given an optimal solution of (35), new values of α generated according to (31) (resulting in most violated constraints) can be added to sets V_{ij}^+ and V_{ij}^- , resulting in tighter relaxations. Note that the use of cuts (as described here) to improve the continuous relaxations of mixed-integer optimization problems is one of the main reasons of the dramatic improvements of MIO software (Bixby, 2012).

In relaxation (35), V_{ij}^+ and V_{ij}^- can be initialized with any (possibly empty) subsets of \mathbb{R}_+ . However, setting $V_{ij}^+ = V_{ij}^- = \{1\}$ yields a relaxation with a simple interpretation, discussed next.

4.3 Diagonally Dominant Matrix Relaxation

Let $\mathbf{\Lambda} \in \mathcal{S}_+^p$ be a diagonally-dominant matrix. Observe that for any $(\mathbf{z}, \boldsymbol{\beta}) \in \{0, 1\}^p \times \mathbb{R}^p$ such that $\boldsymbol{\beta} \circ (\mathbf{e} - \mathbf{z}) = \mathbf{0}$,

$$\begin{aligned} t \geq \boldsymbol{\beta}^\top \mathbf{\Lambda} \boldsymbol{\beta} & \Leftrightarrow t \geq \sum_{i=1}^p \left(\Lambda_{ii} - \sum_{j \neq i} |\Lambda_{ij}| \right) \beta_i^2 + \sum_{i=1}^p \sum_{j=i+1}^p |\Lambda_{ij}| (\beta_i + \text{sign}(\Lambda_{ij})\beta_j)^2 \\ & \Leftrightarrow t \geq \sum_{i=1}^p \left(\Lambda_{ii} - \sum_{j \neq i} |\Lambda_{ij}| \right) \frac{\beta_i^2}{z_i} + \sum_{i=1}^p \sum_{j=i+1}^p |\Lambda_{ij}| \frac{(\beta_i + \text{sign}(\Lambda_{ij})\beta_j)^2}{\min\{1, z_i + z_j\}}, \end{aligned} \quad (36)$$

where the last line follows from using perspective strengthening for the separable quadratic terms and using (7) for the non-separable, rank-one terms. Atamtürk et al. (2021) use a similar strengthening for signal estimation based on nonnegative pairwise quadratic terms.

We now consider using decompositions of the form $\mathbf{\Lambda} + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$, where $\mathbf{\Lambda}$ is a diagonally dominant matrix and $\mathbf{R} \in \mathcal{S}_+^P$. Given such a decomposition, inequalities (36) can be used to strengthen the formulations. Specifically, we consider relaxations of (3) of the form

$$\mathbf{y}^\top \mathbf{y} + \min - 2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \hat{\phi}(\mathbf{z}, \boldsymbol{\beta}) \quad (37a)$$

$$(17b), (17c), (17d), \quad (37b)$$

where

$$\hat{\phi}(\mathbf{z}, \boldsymbol{\beta}) := \max_{\mathbf{\Lambda}, \mathbf{R}} \boldsymbol{\beta}^\top \mathbf{R}\boldsymbol{\beta} + \sum_{i=1}^p \left(\Lambda_{ii} - \sum_{j \neq i} |\Lambda_{ij}| \right) \frac{\beta_i^2}{z_i} + \sum_{i=1}^p \sum_{j=i+1}^p |\Lambda_{ij}| \frac{(\beta_i + \text{sign}(\Lambda_{ij})\beta_j)^2}{\min\{1, z_i + z_j\}} \quad (38a)$$

$$\text{s.t. } \mathbf{\Lambda} + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \quad (38b)$$

$$\Lambda_{ii} \geq \sum_{j < i} |\Lambda_{ji}| + \sum_{j > i} |\Lambda_{ij}| \quad \forall i \in P \quad (38c)$$

$$\mathbf{R} \in \mathcal{S}_+^P. \quad (38d)$$

Proposition 10 *Problem (37) is equivalent to*

$$\mathbf{y}^\top \mathbf{y} + \min - 2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (39a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{z} \leq k \quad (39b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, \quad -\boldsymbol{\beta} \leq \mathbf{u} \quad (39c)$$

$$z_i B_{ii} \geq \beta_i^2 \quad \forall i \in P \quad (39d)$$

$$(sdp_{dd}) \quad 0 \leq w_{ij} \leq 1, \quad w_{ij} \leq z_i + z_j \quad \forall i \neq j \quad (39e)$$

$$0 \geq \frac{\beta_i^2 + 2\beta_i\beta_j + \beta_j^2}{w_{ij}} - 2B_{ij} - B_{ii} - B_{jj} \quad \forall i \neq j \quad (39f)$$

$$0 \geq \frac{\beta_i^2 - 2\beta_i\beta_j + \beta_j^2}{w_{ij}} + 2B_{ij} - B_{ii} - B_{jj} \quad \forall i \neq j \quad (39g)$$

$$\mathbf{B} - \boldsymbol{\beta}\boldsymbol{\beta}' \in \mathcal{S}_+^P \quad (39h)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \quad \mathbf{z} \in [0, 1]^p, \quad \mathbf{u} \in \mathbb{R}_+^p, \quad \mathbf{B} \in \mathbb{R}^{p \times p}. \quad (39i)$$

Proof Let $\mathbf{\Gamma}, \mathbf{\Gamma}^+, \mathbf{\Gamma}^-$ be nonnegative $p \times p$ matrices such that: $\Gamma_{ii} = \Lambda_{ii}$ and $\Gamma_{ij} = 0$ for $i \neq j$; $\Gamma_{ii}^+ = \Gamma_{ii}^- = 0$ and $\Gamma_{ij}^+ - \Gamma_{ij}^- = \Lambda_{ij}$ for $i \neq j$. Problem (38) can be written as

$$\hat{\phi}(\mathbf{z}, \boldsymbol{\beta}) := \max_{\mathbf{\Gamma}, \mathbf{\Gamma}^+, \mathbf{\Gamma}^-} \boldsymbol{\beta}^\top \mathbf{R}\boldsymbol{\beta} + \sum_{i=1}^p \left(\Gamma_{ii} - \sum_{j \neq i} (\Gamma_{ij}^+ + \Gamma_{ij}^-) \right) \frac{\beta_i^2}{z_i} \quad (40a)$$

$$+ \sum_{i=1}^p \sum_{j=i+1}^p \left(\Gamma_{ij}^+ \frac{(\beta_i + \beta_j)^2}{\min\{1, z_i + z_j\}} + \Gamma_{ij}^- \frac{(\beta_i - \beta_j)^2}{\min\{1, z_i + z_j\}} \right) \quad (40b)$$

$$\text{s.t. } \mathbf{\Gamma} + \mathbf{\Gamma}^+ + \mathbf{\Gamma}^- + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \quad (40c)$$

$$\Gamma_{ii} \geq \sum_{j < i} (\Gamma_{ji}^+ + \Gamma_{ji}^-) + \sum_{j > i} (\Gamma_{ij}^+ + \Gamma_{ij}^-) \quad \forall i \in P \quad (40d)$$

$$\mathbf{R} \in \mathcal{S}_+^P. \quad (40e)$$

Then, similarly to the proof of Theorem 6, it is easy to show that the dual of (40) is precisely (39). \blacksquare

4.4 Relaxing the $(p + 1)$ -Dimensional PSD Constraint

We now discuss a relaxation of the p -dimensional semidefinite constraint $\mathbf{B} - \boldsymbol{\beta}\boldsymbol{\beta}^\top \in \mathcal{S}_+^p$, present in all formulations. Let \mathbf{V} be a matrix whose j -th column \mathbf{V}_j is an eigenvector of $\mathbf{X}^\top \mathbf{X}$. Consider the optimization problem

$$\underline{\phi}_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta}) := \max_{\mathbf{A}_T, \mathbf{R}, \boldsymbol{\pi}} \boldsymbol{\beta}^\top \mathbf{R} \boldsymbol{\beta} + \sum_{T \in \mathcal{P}} \frac{\boldsymbol{\beta}_T^\top \mathbf{A}_T \boldsymbol{\beta}_T}{\min\{1, z(T)\}} \quad (41a)$$

$$\text{s.t. } \sum_{T \in \mathcal{P}} \mathbf{A}_T + \mathbf{R} = \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I} \quad (41b)$$

$$\mathbf{A}_T \in \mathcal{S}_+^T \quad \forall T \in \mathcal{P} \quad (41c)$$

$$\mathbf{R} = \mathbf{V} \text{diag}(\boldsymbol{\pi}) \mathbf{V}^\top \quad (41d)$$

$$\boldsymbol{\pi} \in \mathbb{R}_+^n. \quad (41e)$$

Observe that the objective and constraints (41a)–(41c) are identical to (15). However, instead of (15e), we have $\mathbf{R} = \sum_{j=1}^{\min\{p, n\}} \pi_j \mathbf{V}_j \mathbf{V}_j^\top$. Moreover, since $\boldsymbol{\pi} \geq \mathbf{0}$, $\mathbf{R} \in \mathcal{S}_+^p$ in any feasible solution of (41), thus (15) is a relaxation of (41), and, hence, $\underline{\phi}_{\mathcal{P}}$ is indeed a lower bound on $\phi_{\mathcal{P}}$. Finally, (41) is feasible if $\lambda = 0$ or \mathcal{P} contains all singletons, as it is possible to set $\mathbf{A}_{\{i\}} = \lambda$, $\mathbf{A}_T = 0$ for $|T| > 1$, and set $\boldsymbol{\pi}$ equal to the eigenvalues of $\mathbf{X}^\top \mathbf{X}$. Therefore, instead of (17), one may use the simpler convex relaxation

$$\mathbf{y}^\top \mathbf{y} + \min - 2\mathbf{y}^\top \mathbf{X} \boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \underline{\phi}_{\mathcal{P}}(\mathbf{z}, \boldsymbol{\beta}) \quad (42a)$$

$$\mathbf{e}^\top \mathbf{z} \leq k \quad (42b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, \quad -\boldsymbol{\beta} \leq \mathbf{u} \quad (42c)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \quad \mathbf{z} \in [0, 1]^p, \quad \mathbf{u} \in \mathbb{R}_+^p \quad (42d)$$

for (1).

Proposition 11 *If $\mathcal{P} = \{T \subseteq P : |T| \leq 2\}$, then problem (42) is equivalent to*

$$\mathbf{y}^\top \mathbf{y} + \min -2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \mathbf{e}^\top \mathbf{u} + \langle \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}, \mathbf{B} \rangle \quad (43a)$$

$$s.t. \mathbf{e}^\top \mathbf{z} \leq k \quad (43b)$$

$$\boldsymbol{\beta} \leq \mathbf{u}, -\boldsymbol{\beta} \leq \mathbf{u} \quad (43c)$$

$$\begin{pmatrix} z_i & \beta_i \\ \beta_i & B_{ii} \end{pmatrix} \succeq 0 \quad \forall i = 1, \dots, p \quad (43d)$$

$$(sdp_{LB}) \quad 0 \leq w_{ij} \leq \min\{1, z_i + z_j\} \quad \forall i < j \quad (43e)$$

$$\begin{pmatrix} w_{ij} & \beta_i & \beta_j \\ \beta_i & B_{ii} & B_{ij} \\ \beta_j & B_{ij} & B_{jj} \end{pmatrix} \succeq 0 \quad \forall i < j \quad (43f)$$

$$\mathbf{V}_j^\top (\mathbf{B} - \boldsymbol{\beta}\boldsymbol{\beta}^\top) \mathbf{V}_j \geq 0 \quad \forall j = 1, \dots, \min\{n, p\} \quad (43g)$$

$$\boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{z} \in [0, 1]^p, \mathbf{u} \in \mathbb{R}_+^p, \mathbf{B} \in \mathbb{R}^{p \times p}. \quad (43h)$$

Proof The proof is based on conic duality similar to the proof of Theorem 6. ■

Observe that in formulation (43), the $(p + 1)$ -dimensional semidefinite constraint (19f) is replaced with $\min\{p, n\}$ rank-one quadratic constraints (43g). We denote by sdp_{LB} the relaxation of sdp_2 obtained by replacing (21f) with (43g). In general, sdp_{LB} is still an SDP due to constraints (43f); however, note that sdp_{LB} can be implemented in a conic quadratic framework by using cuts, as described in §4.2. Moreover, constraints (43g) could also be dynamically refined to better approximate the SDP constraint, or formulation (43) could be improved with ongoing research on approximating SDP via mixed-integer conic quadratic optimization (Kocuk et al., 2016, 2018).

Remark 12 *We observe that formulation (43) is solved substantially faster than sdp_2 (with Mosek) with constraints (43f) formulated as semi-definite constraints. Indeed, the $\mathcal{O}(p^2)$ low-dimensional constraints (22f) can actually be handled efficiently, but the major computational bottleneck towards solving sdp_2 is handling the single large-dimensional positive semi-definite constraint (22g).*

Remark 13 *Since the first submission of this paper, additional relaxations and approximations of sdp_r have been proposed in the literature, see Ben-Ameur (2024).*

5 Computations

In this section, we report computational experiments with the proposed conic relaxations on synthetic as well as benchmark datasets. Semidefinite optimization problems are solved with MOSEK 8.1 solver, and conic quadratic optimization problems (continuous and mixed-integer) are solved with CPLEX 12.8 solver. All computations are performed on a laptop with a 1.80GHz Intel®Core™ i7-8550U CPU and 16 GB main memory. All solver parameters were set to their default values. We divide our discussion in two parts: first, in §5.2, we focus on the relaxation quality of sdp_r and its ability to approximate the exact ℓ_0 -problem (1); then, in §5.3, we adopt the same experimental framework used by Bertsimas et al.

(2016); Hastie et al. (2017) to generate synthetic instances and evaluate the proposed conic formulations from an inference perspective. In both cases, our results compare favorably with existing approaches in the literature. Finally, in §5.4, we summarize our findings and discuss possible extensions.

5.1 Datasets

We use the benchmark datasets in Table 1. The first five were first used by Miyashiro and Takano (2015) in the context of MIO algorithms for best subset selection, and later used by Gómez and Prokopyev (2021). The `diabetes` dataset with all second interactions was introduced by Efron et al. (2004) in the context of `lasso`, and later used by Bertsimas et al. (2016). A few datasets require some manipulation to eliminate missing values and handle categorical variables. The processed datasets before standardization⁵ can be downloaded from <http://atamturk.ieor.berkeley.edu/data/sparse.regression>.

In addition, we also use synthetic datasets generated similarly to those used by Bertsimas et al. (2016); Hastie et al. (2017). Here we present a summary of the simulation setup and refer the readers to Hastie et al. (2017) for an extended description. For given dimensions n, p , sparsity s , predictor autocorrelation ρ , and signal-to-noise ratio SNR, the instances are generated as follows:

1. The (true) coefficients β_0 have the first s components equal to one, and the rest equal to zero.
2. The rows of the predictor matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ are drawn from i.i.d. distributions $\mathcal{N}_p(\mathbf{0}, \Sigma)$, where $\Sigma \in \mathbb{R}^{p \times p}$ has entry (i, j) equal to $\rho^{|i-j|}$.
3. The response vector $\mathbf{y} \in \mathbb{R}^n$ is drawn from $\mathcal{N}_p(\mathbf{X}\beta_0, \sigma^2\mathbf{I})$, where $\sigma^2 = \beta_0^\top \mathbf{X}\beta_0 / \text{SNR}$.

Similar data generation has been used in the literature (Bertsimas et al., 2016; Hastie et al., 2017).

5.2 Relaxation Quality

In this section, we test the ability of `sdp2`, given in (22), and of `sdpLB`, given in (43), to provide near-optimal solutions to problem (1), and compare its performance with MIO approaches. In §5.2.1, we focus on the pure best subset selection problem with $\lambda = 0$, which has received relatively little attention in the literature (Bertsimas et al., 2016); in §5.2.2 we consider problems with ℓ_0 - ℓ_2 regularization, which has received more attention in the literature (Bertsimas and Van Parys, 2020; Hazimeh and Mazumder, 2020; Hazimeh et al., 2022; Xie and Deng, 2020); in §5.2.3 we study the impact of model complexity parameter r on the relaxation quality, and in §5.2.4 we study the scalability of the proposed methods.

Computing optimality gaps for `sdpr` The optimal objective value ν_ℓ^* of `sdpr` provides a lower bound on the optimal objective value of (1). To obtain an upper bound, we use a simple greedy heuristic to retrieve a feasible solution for (1): given an optimal solution vector $\bar{\beta}^*$ for `sdpr`, let $\bar{\beta}_{(k)}^*$ denote the k -th largest absolute value. For $T = \left\{ i \in P : |\bar{\beta}_i^*| \geq \bar{\beta}_{(k)}^* \right\}$,

5. In our experiments, the datasets were standardized first.

let $\hat{\beta}_T$ be the k -dimensional `ols/ridge` estimator using only predictors in T , i.e.,

$$\hat{\beta}_T = (\mathbf{X}_T^\top \mathbf{X}_T + \lambda \mathbf{I}_T)^{-1} \mathbf{X}_T^\top \mathbf{y},$$

where \mathbf{X}_T denotes the $n \times k$ matrix obtained by removing the columns with indexes not in T , and let $\tilde{\beta}$ be the P -dimensional vector obtained by filling the missing entries in $\hat{\beta}_T$ with zeros. Since $\|\tilde{\beta}\|_0 \leq k$ by construction, $\tilde{\beta}$ is feasible for (1), and its objective value ν_u is an upper bound on the optimal objective value of (1). Moreover, the optimality gap provided by any approach can be computed as

$$\text{gap} = \frac{\nu_u - \nu_\ell^*}{\nu_\ell^*} \times 100. \tag{44}$$

While stronger relaxations result in improved lower bounds ν_ℓ^* , the corresponding heuristic upper bounds ν_u are not necessarily better; thus, the optimality gaps are not guaranteed to improve with stronger relaxations. Nevertheless, as shown next, stronger relaxations, in general, yield much smaller gaps in practice.

We point out that the main focus of the strong relaxations is to obtain improved lower bounds ν_ℓ^* . Randomized rounding methods (Pilanci et al., 2015; Xie and Deng, 2020), more sophisticated rounding heuristics (Dong et al., 2015), or alternative heuristic methods (Hazimeh and Mazumder, 2020) can be used to obtain improved upper bounds. Nevertheless, the quality of the upper bounds obtained from the greedy rounding method can be used to estimate how well the solutions from the relaxations match the sparsity pattern of the optimal solution.

5.2.1 $\lambda = 0$ CASE

For each dataset with $\lambda = \mu = 0$, we solve the conic relaxations of (1) `sdp1` and `sdp2` as well as `sdpLB` and the mixed-integer formulation `big-M` given by (3)–(4). In our experiments, we set $M = 3\|\beta_{\text{ols}}\|_\infty$, where β_{ols} is the ordinary least square estimator⁶ and set a time limit of 10 minutes. For data with $p \leq 40$ we solve problems with cardinalities $k \in \{3, \dots, 10\}$, and for `diabetes` and `crime` we solve problems with $k \in \{3, \dots, 30\}$. Table 2 shows, for each dataset and method, the average lower bound (LB) and upper bound (UB) found by each method, the `gap` (44), and the `time` required to solve the problems (in seconds) – the average is taken across all k values. In all cases, lower and upper bounds are scaled so that the best upper bound for any given instance has value $\nu_u^* = 100$.

The `big-M` method is highly inconsistent and prone to numerical difficulties due to the use of big- M constraints. First, for three datasets (`servo`, `auto MPG` and `breast cancer`) the method fails due to numerical issues (“failure to solve MIP subproblem”). In addition, for `solar flare` the solver reports very fast solution times but the solutions are, in fact, infeasible for problem (1): by default in CPLEX, if $z_i \leq 10^{-5}$ in a solution then z_i is deemed to satisfy the integrality constraint $z_i \in \{0, 1\}$. Thus, if the big- M constant is large enough, then constraint (4) may, in fact, allow nonzero values for β_i even when “ $z_i = 0$ ”. In particular, in `solar flare` we found that the solution β_{mio} reported by the MIO solver satisfies⁷ $\|\beta_{\text{mio}}\|_0 = 20$, regardless of the value of k used, violating the sparsity constraint.

6. Bertsimas et al. (2016) set $M = 2\|\hat{\beta}\|_\infty$ for some heuristic solution $\hat{\beta}$

7. We consider $\beta_i \neq 0$ whenever $\|\beta_i\| > 10^{-4}$.

dataset	method	LB	UB	gap(%)	time
housing	sdp ₁	99.4±0.6	100.1±0.1	0.7±0.0	0.03±0.02
	sdp ₂	99.6±0.6	100.1±0.1	0.5±0.6	0.07±0.03
	sdp _{LB}	98.8±0.6	100.4±0.1	1.6±0.8	0.06±0.03
	big-M	100.0±0.0	100.0±0.0	0.0±0.0	0.01±0.01
servo	sdp ₁	86.8±5.5	109.5±10.3	27.3±20.6	0.02±0.01
	sdp ₂	94.9±2.9	106.2±16.5	12.2±19.8	0.10±0.01
	sdp _{LB}	89.5±2.7	109.2±15.5	21.8±14.9	0.17±0.03
	big-M†	†	†	†	†
auto MPG	sdp ₁	75.3±10.3	115.3±6.0	55.8±23.7	0.07±0.04
	sdp ₂	96.7±3.3	100.5±0.8	4.0±4.2	0.24±0.02
	sdp _{LB}	78.8±7.7	101.6±2.7	30.0±14.0	0.40±0.09
	big-M†	†	†	†	†
solar flare	sdp ₁	97.5±1.5	103.3±1.1	6.0±2.0	0.07±0.03
	sdp ₂	99.2±0.8	100.0±0.0	1.0±0.6	0.28±0.06
	sdp _{LB}	97.8±1.6	102.3±1.9	4.6±2.7	0.13±0.02
	big-M††	98.1±1.7	98.1±1.7	-	0.01±0.01
breast cancer	sdp ₁	88.9±3.1	101.5±1.7	14.4±5.6	0.15±0.02
	sdp ₂	98.0±0.6	100.4±0.8	2.4±1.1	0.77±0.07
	sdp _{LB}	94.8±0.5	100.5±0.7	6.0±0.5	0.40±0.03
	big-M†	†	†	†	†
diabetes	sdp ₁	95.2±3.2	115.2±11.8	22.2±16.3	3.58±0.77
	sdp ₂	97.4±1.3	105.4±4.2	8.2±5.2	9.28±1.12
	sdp _{LB} †	†	†	†	†
	big-M	99.0±0.9	100.0±0.0	1.0±0.9	416.17±260.57
crime	sdp ₁	97.8±1.3	103.2±2.4	5.6±3.6	17.82±0.98
	sdp ₂	99.0±0.8	101.6±2.0	2.7±2.7	45.29±4.06
	sdp _{LB}	94.6±2.0	109.7±2.8	16.0±4.9	5.87±0.43
	big-M	96.4±1.7	100.0±0.0	3.7±1.8	527.03±185.64

† Error in solving problem.

†† Infeasible solution is reported as optimal.

Table 2: Results with $\lambda = 0$ on real instances. Lower and upper bounds are scaled so that the best upper bound for a given instance is 100. Mean \pm stdev are reported.

We also point out that sdp_{LB} struggles with numerical difficulties in **diabetes**: the problems

are incorrectly found to be unbounded. In contrast, sdp_r methods are solved without numerical difficulties.

In terms of the relaxation quality, we find that sdp_2 is the best as expected. It consistently delivers better lower and upper bounds compared to the other conic relaxations and even outperforms big-M in terms of lower bounds and gaps in the largest dataset (*crime*). The strength of the relaxation comes at the expense 2–4-fold larger computation time than sdp_1 , but on the other hand sdp_2 is substantially faster than big-M on large datasets. We see that neither sdp_1 nor sdp_{LB} dominates each other in terms of relaxation quality. While sdp_1 is faster on the smaller datasets, sdp_{LB} is faster on *crime*, indicating that sdp_{LB} may scale better (we corroborate this statement in §5.2.4). Finally, big-M , in datasets where numerical issues do not occur, is able to find high-quality solutions consistently but struggles to find matching lower bound in larger instances despite significantly higher computation time spent.

Figures 4 and 5 present detailed results on lower bounds and gaps as a function of the sparsity parameter k for the *diabetes* and *crime* datasets. For small values of k , big-M is arguably the best method, solving the problems to optimality. However, as k increases, the quality of the lower bounds and gaps deteriorate: for *diabetes*, sdp_2 finds better solutions than big-M for $k \geq 18$; for *crime*, sdp_1 and sdp_2 find better lower bounds for $k \geq 8$ (and, in the case of sdp_2 , better gaps as well), and sdp_{LB} matches the lower bound found by big-M for $k \geq 14$, despite requiring only five seconds (instead of 10 minutes) to find such lower bounds. Observe that the number of possible supports $\binom{p}{k} = \mathcal{O}(p^k)$ for problem (1) scales exponentially with k , thus enumerative methods such as branch-and-bound may struggle as k grows.

5.2.2 $\lambda > 0$ CASE

For each dataset with⁸ $\lambda = 0.05$ and $\mu = 0$, we solve the conic relaxations of (1) sdp_1 , sdp_2 and sdp_{LB} and the “big- M free” mixed-integer formulation (5) with a time limit of

8. Since data is standardized so that each column has a unit norm, a value of $\lambda = 0.05$ corresponds to an increase of 5% in the diagonal elements of the matrix $\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I}$.

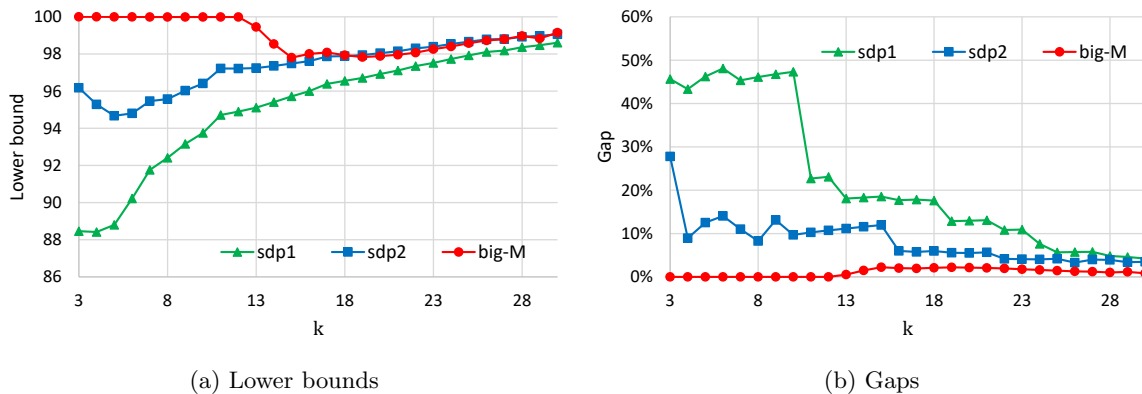



Figure 4: Detailed results on the *diabetes* dataset with $\lambda = 0$.

Figure 5: Detailed results on the `crime` dataset with $\lambda = 0$.

10 minutes (`persp`). This MIO formulation is possible since $\lambda > 0$, and has been shown to be competitive (Xie and Deng, 2020; Hazimeh et al., 2022) with the tailored algorithm proposed by Bertsimas and Van Parys (2020). For datasets with $p \leq 40$ we solve the problems with cardinalities $k \in \{3, \dots, 10\}$, and for `diabetes` and `crime` we solve the problems with $k \in \{3, \dots, 30\}$. Table 3 shows, for each dataset and method, the average lower bound (LB) and upper bound (UB) found by each method, the `gap` (44), and the `time` required to solve the problems (in seconds) – the average is taken across all k values. In all cases, lower and upper bounds are scaled so that the best upper bound for any given instance has value $\nu_u^* = 100$.

We observe that instances with $\lambda = 0.05$ are much easier to solve than those with $\lambda = 0$: no numerical issues occur for `sdpLB` or `persp`, and lower and upper bounds are much better for all methods. The mixed integer formulation `persp` comfortably solves the small instances with $p \leq 40$ to optimality, but `sdp2` yields better lower bounds and gaps for the larger instances `diabetes` and `crime` in a fraction of the time used by `persp`.

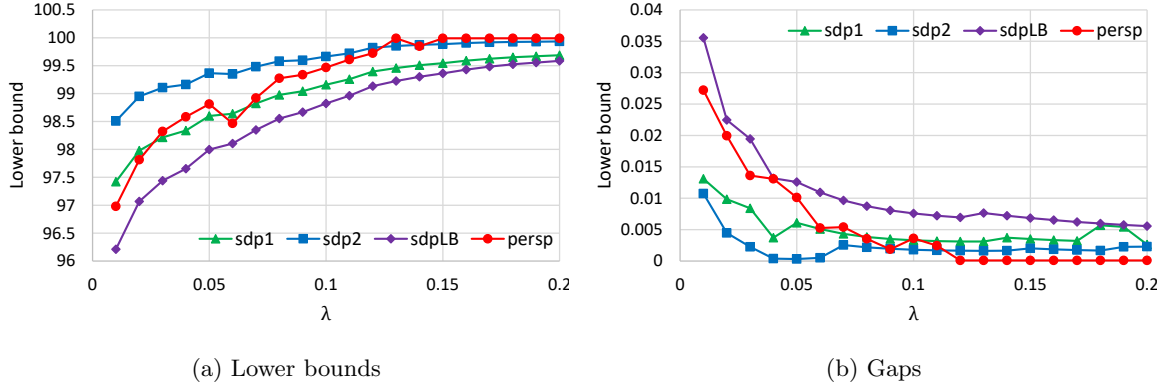
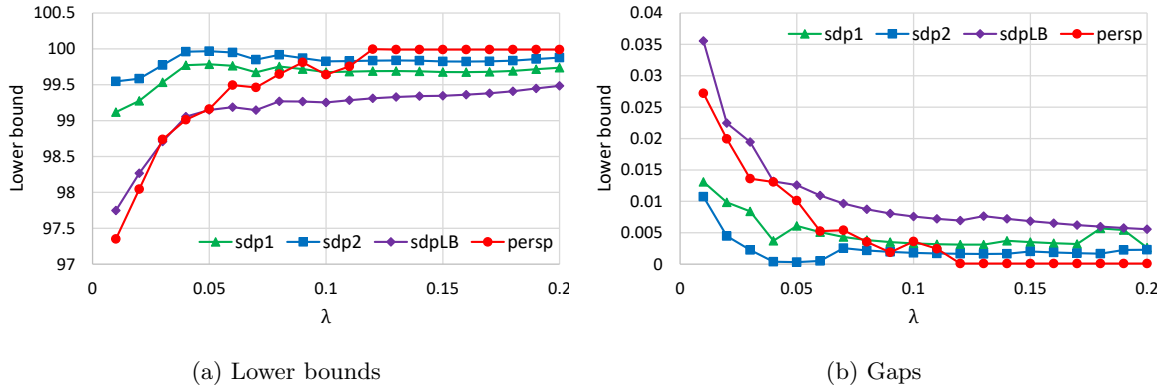
Figures 6 and 7 present lower bounds and gaps as a function of the regularization parameter λ , for `diabetes` and `crime` datasets (with $k = 15$). We observe that for a low value of λ , `persp` struggles to find good lower bounds, e.g., it is outperformed by all conic relaxations in `crime` for $\lambda \leq 0.02$, and is worse than `sdp2` for $\lambda \leq 0.1$ in terms of lower bounds and gaps in both datasets. As λ increases, all methods deliver better bounds, and `persp` is eventually able to solve all problems to optimality.

As expected, the performance of `persp` improves as λ increases. The perspective relaxation discussed in §1.3 exploits the separable terms introduced by the ℓ_2 -regularization: as λ increases, these separable terms have a larger weight in the objective, and the strength of the relaxation improves as a consequence. Note that the conic relaxations also improve with larger λ : they are based on decompositions of the matrix $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ into one- and two-variable terms, and the addition of the separable terms allows for a much richer set of decompositions. For large values of λ , $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ becomes highly diagonal dominant, and the perspective relaxation alone provides a substantial strengthening. In this case, the advanced conic relaxations have a marginal impact and MIO methods with perspective strengthening performs better overall. In contrast, for low values of λ , the conic relaxations

dataset	method	LB	UB	gap(%)	time
housing	sdp ₁	99.7±0.4	100.2±0.3	0.5±0.6	0.03±0.02
	sdp ₂	99.8±0.3	100.1±0.2	0.3±0.5	0.06±0.02
	sdp _{LB}	99.5±0.4	100.3±0.3	0.8±0.6	0.06±0.02
	persp	100.0±0.0	100.0±0.0	0.0±0.0	0.11±0.03
servo	sdp ₁	95.9±3.0	102.2±6.7	6.7±4.1	0.03±0.01
	sdp ₂	99.5±0.5	100.6±1.1	1.1±1.6	0.11±0.01
	sdp _{LB}	97.6±1.4	102.0±2.1	4.6±3.3	0.16±0.02
	persp	100.0±0.0	100.0±0.0	0.0±0.0	0.28±0.13
auto MPG	sdp ₁	89.1±6.1	101.4±1.2	14.4±8.5	0.05±0.01
	sdp ₂	99.8±0.2	100.0±0.1	0.2±0.3	0.25±0.04
	sdp _{LB}	92.7±3.1	101.1±1.5	9.2±4.0	0.35±0.02
	persp	100.0±0.0	100.0±0.0	0.0±0.0	1.29±0.60
solar flare	sdp ₁	99.3±0.5	100.1±0.1	0.8±0.5	0.07±0.01
	sdp ₂	99.9±0.1	100.1±0.1	0.2±0.1	0.28±0.03
	sdp _{LB}	99.2±0.7	100.4±1.2	1.2±1.0	0.16±0.03
	persp	100.0±0.0	100.0±0.0	0.0±0.0	1.75±1.07
breast cancer	sdp ₁	94.9±1.8	100.8±0.4	6.3±2.4	0.18±0.04
	sdp ₂	99.6±0.2	100.1±0.2	0.5±0.3	0.72±0.06
	sdp _{LB}	97.5±0.6	100.5±0.4	2.9±0.9	0.36±0.05
	persp	100.0±0.0	100.0±0.0	0.0±0.0	56.12±44.34
diabetes	sdp ₁	98.9±0.6	100.2±0.2	1.2±0.7	2.13±0.24
	sdp ₂	99.6±0.2	100.1±0.1	0.5±0.3	5.83±0.79
	sdp _{LB}	98.2±1.3	100.3±0.3	2.2±1.4	1.48±0.18
	persp	99.4±0.5	100.0±0.0	0.6±0.5	441.90±258.29
crime	sdp ₁	99.3±0.9	100.3±0.9	1.1±1.7	19.15±1.30
	sdp ₂	99.7±0.4	100.2±0.8	0.5±1.0	43.86±2.38
	sdp _{LB}	98.7±1.0	100.7±1.3	2.0±2.3	5.30±0.35
	persp	99.5±0.4	100.1±0.1	0.6±0.4	518.03±175.65

Table 3: Results with $\lambda = 0.05$ on real instances. Lower and upper bounds are scaled so that the best upper bound for a given instance is 100. Mean \pm stdev are reported.

result in substantial strengthening over the perspective relaxation, and sdp_r outperforms persp as a consequence.


 Figure 6: Detailed results on the `diabetes` dataset with $k = 15$.

 Figure 7: Detailed results on the `crime` dataset with $k = 15$.

5.2.3 THE EFFECT OF MODEL COMPLEXITY r

In §5.2.1–5.2.2 we reported computations with sdp_r with $r \leq 2$. In experiments with those datasets, sdp_3 yields almost the same strengthening as sdp_2 , but with much larger computational cost. Since sdp_2 already achieves gaps close to 0 in those instances, there is little room for improvement with higher values of r .

If the matrix $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ has high diagonal dominance, which happens if $n > p$ or if λ is large, then there are many ways to decompose it into low-dimensional rank-one terms, and sdp_r with r small achieves good relaxations. In contrast, if the matrix $\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}$ has low diagonal dominance, it may be difficult to extract low-dimensional rank-one terms. In the extreme case of a rank-one case matrix, while sdp_p results in the convex description, sdp_r with $r < p$ achieves no improvement. In this section, we illustrate this phenomenon.

First, we test on small synthetic instances with $p = 15$ and $n = 10$, setting the true sparsity to $s = 5$, autocorrelation $\rho = 0.35$, signal-noise-ratio $\text{SNR} \in \{1, 5\}$, sparsity $k \in \{3, 4, 5, 6, 7, 8\}$, and for each combination of parameters we generate five instances. We report in Figure 8 the gaps obtained by sdp_r for different values of r and λ – averaging across instances and different values of SNR and k . In addition, Figure 9 depicts the distribution of

computational times required to solve the problems. We observe that for $\lambda = 0$, sdp_r with $r \leq 4$ results in no strengthening and gaps of 100%; sdp_5 results in a small improvement (note that $5 = p - n$), while sdp_r with $r \geq 6$ results in larger improvements. These results suggest that, with $\lambda = 0$, stronger formulations require rank-one strengthening with at least $p - n$ variables. We also observe that, as λ increases, the gaps reported by all methods decrease substantially, and the incremental strengthening obtained from larger values of r decreases: for $\lambda \geq 0.05$ sdp_4 performs almost identical to sdp_8 , and for $\lambda = 0.15$ sdp_3 is similar to sdp_8 and sdp_2 already results in low optimality gaps. The computational time required to solve sdp_r scales exponentially with r since the number of constraints increases exponentially as well. We conclude that sdp_2 is well suited for the $n > p$ case or for medium values of λ (for larger values sdp_1 or even the simple perspective relaxation may be preferable), while sdp_r with $r \geq 3$ achieves a good improvement in relaxation quality for low values of λ , at the expense of larger computational times.

Next, we test sdp_r on the **housing** dataset with $\lambda = 0$. Since this dataset has high diagonal dominance (see Table 1), even sdp_1 results in small optimality gaps (see the first row of Table 2). Thus, instead of using all $n = 506$ datapoints, we randomly select only $n_0 < n$ datapoints, resulting in small diagonal dominance – for $n_0 = 10$, matrix $\mathbf{X}^\top \mathbf{X}$ is rank-deficient. Figure 10 shows the optimality gaps as a function of r and n_0 . The relaxation quality improves with increasing n_0 and increasing r . For $n_0 = 20$, setting $r \geq 5$ produces optimality gaps close to 0% while sdp_1 results in a gap of 13%, sdp_2 a gap of 8%, and sdp_3 yields a gap of 3%. For $n_0 = 50$, sdp_2 results in a gap of 2%, and sdp_3 yields a gap almost equal to 0%. Overall, the results are consistent with the experiments on synthetic data.

5.2.4 ON SCALABILITY

As discussed in §5.2.3, formulation sdp_r for large values of r can be expensive to solve. Moreover, even sdp_1 and sdp_2 are semidefinite programs, which may not scale well for large

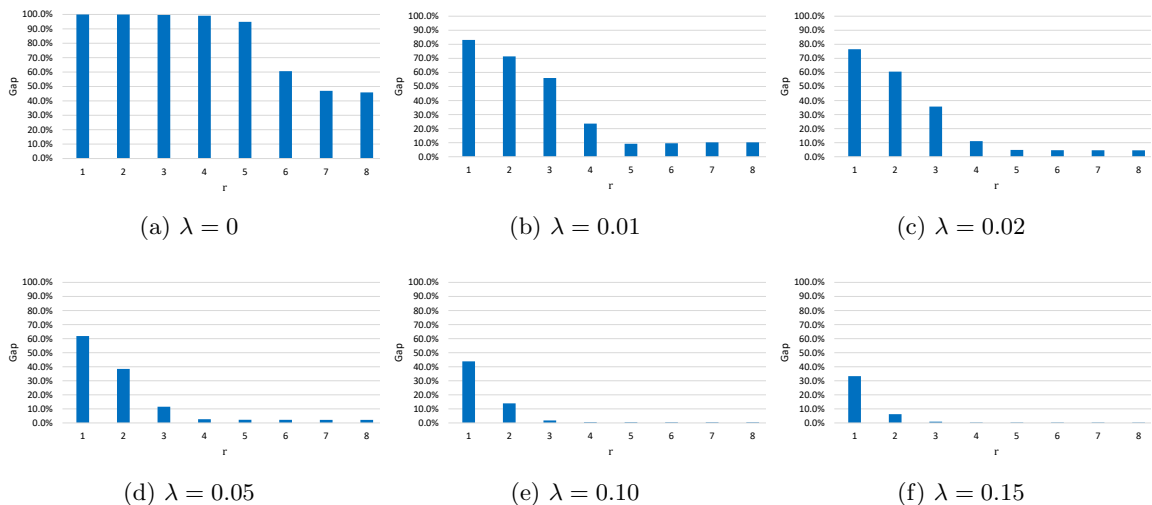




Figure 8: Optimality gaps of sdp_r with synthetic data, $1 \leq r \leq 8$.


 Figure 9: Time required to solve sdp_r , $1 \leq r \leq 8$.

 Figure 10: Optimality gaps of sdp_r with the altered housing dataset (using only n_0 datapoints), $1 \leq r \leq 8$.

values of p . In this section, we present computations illustrating that while this is indeed the case, formulation sdp_{LB} —which replaces the semidefinite constraint $\mathbf{B} - \beta\beta' \in S_+^P$ with the quadratic constraints (43g)—scales much better and in fact can significantly outperform persp in terms of relaxation quality.

We generate synthetic instances with $p \in \{100, 150, \dots, 500\}$, $n = 500$, true sparsity parameter $s = 30$, autocorrelation $\rho = 0.35$, signal-noise-ratio $\text{SNR} \in \{1, 5\}$, sparsity $k = 30$; for each combination of parameters we generate five instances, and solve them for $\lambda \in \{0.01, 0.02, 0.05, 0.15\}$ and $\mu = 0$. Table 4 reports, for sdp_1 , sdp_2 , sdp_{LB} and persp —using formulation (5) with a time limit of 600 seconds—, the time required to solve the problems and the optimality gap proven.

We observe that persp is unable to solve the problems within the ten-minute time limit and results in larger gaps than all other approaches despite using substantially more time in most cases. We also observe that sdp_r formulations struggle in instances with $p \geq 200$. Interestingly, sdp_2 requires consistently 2-4 times more than sdp_1 regardless of the dimension p . A similar factor was observed in Tables 2 and 3 with real data, suggesting that computational times with sdp_2 are within the same order-of-magnitude as sdp_1 . Finally, sdp_{LB} is substantially faster than both sdp_1 and sdp_2 . While it results in larger gaps than

p	<u>sdp₁</u>		<u>sdp₂</u>		<u>sdp_{LB}</u>		<u>persp</u>	
	time(s)	gap(%)	time(s)	gap(%)	time(s)	gap(%)	time(s)	gap(%)
100	19±3	0.5±0.6	44±9	0.1±0.1	5±1	1.0±1.1	TL	3.3±4.2
150	153±19	1.1±1.5	356±61	0.2±0.4	20±2	1.9±2.2	TL	6.2±7.2
200	673±64	2.7±3.0	1,691±165	0.6±0.9	42±3	4.3±4.0	TL	12.5±10.5
250	†	†	†	†	79±4	7.1±5.6	TL	17.2±13.2
300	†	†	†	†	147±7	12.2±7.6	TL	21.7±14.6
350	†	†	†	†	248±14	17.6±11.0	TL	25.9±17.0
400	†	†	†	†	391±36	24.0±14.9	TL	29.1±18.9
450	†	†	†	†	394±46	32.2±18.3	TL	34.3±21.0
500	†	†	†	†	462±43	39.3±21.9	TL	38.8±22.8

Table 4: Computational times and gaps on synthetic instances as a function of p . TL= Time Limit. †= Unable to solve (either due to very large computational times or memory issues). Numbers after “±” are the sample standard deviation.

sdp_2 as expected since the high-dimensional constraint (22g) is relaxed, it still yields better optimality gaps than persp .

5.3 Inference study on synthetic instances

We now present inference results on synthetic data using the same simulation setup as by Bertsimas et al. (2016); Hastie et al. (2017), see Hastie et al. (2017) for an extended description. Specifically, we generate synthetic data as described in §5.1, and use the evaluation metrics used by Hastie et al. (2017), described next.

5.3.1 EVALUATION METRICS

Let \mathbf{x}_0 denote the test predictor drawn from $\mathcal{N}_p(\mathbf{0}, \Sigma)$ and let y_0 denote its associated response value drawn from $\mathcal{N}(\mathbf{x}_0^\top \beta_0, \sigma^2)$. Given an estimator $\hat{\beta}$ of β_0 , the following metrics are reported:

Relative risk

$$\text{RR}(\hat{\beta}) = \frac{\mathbb{E} \left(\mathbf{x}_0^\top \hat{\beta} - \mathbf{x}_0^\top \beta_0 \right)^2}{\mathbb{E} \left(\mathbf{x}_0^\top \beta_0 \right)^2}$$

with a perfect score 0 and null score of 1.

Relative test error

$$\text{RTE}(\hat{\beta}) = \frac{\mathbb{E} \left(\mathbf{x}_0^\top \hat{\beta} - y_0 \right)^2}{\sigma^2}$$

with a perfect score of 1 and null score of SNR+1.

Proportion of variance explained

$$1 - \frac{\mathbb{E} \left(\mathbf{x}_0^\top \hat{\boldsymbol{\beta}} - y_0 \right)^2}{\text{Var}(y_0)}$$

with perfect score of $\text{SNR}/(1+\text{SNR})$ and null score of 0.

Sparsity We record the number of nonzeros⁹, $\|\hat{\boldsymbol{\beta}}\|_0$, as done by Hastie et al. (2017). Additionally, we also report the number of variables correctly identified, given by $\sum_{i=1}^p \mathbb{1}\{\hat{\beta}_i \neq 0 \text{ and } (\beta_0)_i \neq 0\}$.

5.3.2 PROCEDURES

In addition to the training set of size n , a validation set of size n is generated with the same parameters, matching the precision of leave-one-out cross-validation. We use the following procedures to obtain estimators $\hat{\boldsymbol{\beta}}$.

elastic net We solve the elastic net procedure using the parametrization

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda (\alpha \|\boldsymbol{\beta}\|_1 + (1 - \alpha) \|\boldsymbol{\beta}\|_2^2)$$

where $\alpha, \lambda \geq 0$ are the regularization parameters. We let $\alpha = 0.1\ell$ for integer $0 \leq \ell \leq 10$, we generated 50 values of λ ranging from $\lambda_{max} = \|\mathbf{X}^\top \mathbf{y}\|_\infty$ to $\lambda_{max}/200$ on a log scale, and using the pair (λ, μ) that results in the best prediction error on the validation set. A total of 500 (α, λ) pairs are tested.

sdp₂ The estimator obtained from solving **sdp₂** ($\lambda = \mu = 0$) for all values of $k = 0, \dots, 7$ and choosing the one that results in the best prediction error on the validation set.

The **elastic net** procedure approximately corresponds to the **lasso** procedure with 100 tuning parameters used by Hastie et al. (2017). Similarly, **sdp₂** with cross-validation approximately corresponds to the **best subset** procedure with 51 tuning parameters¹⁰ used by Hastie et al. (2017); nonetheless, the estimators from Hastie et al. (2017) are obtained by running an MIO solver for 3 minutes, while ours are obtained from solving to optimality a strong convex relaxation.

5.3.3 OPTIMALITY GAPS AND COMPUTATION TIMES

Before describing the statistical results, we briefly comment on the relaxation quality and computation time of **sdp₂**. Table 5 shows, for instances with $n = 500$, $p = 100$, and $s = 5$, the optimality gap and relaxation quality of **sdp₂** — each column represents the average over ten instances generated with the same parameters. In all cases, **sdp₂** produces optimal

9. An entry $\hat{\beta}_i$ is deemed to be non-zero if $|\hat{\beta}_i| > 10^{-5}$. This is the default integrality precision in commercial MIO solvers.

10. Hastie et al. (2017) use values of $k = 0, \dots, 50$. Nonetheless, in our computations with the same tuning parameters, we found that values of $k \geq 8$ are never selected after cross-validation. Thus, our procedure with 8 tuning parameters results in the same results as the one with 51 parameters from a statistical viewpoint but requires only a fraction of the computational effort.

SNR		0.05	0.09	0.14	0.25	0.42	0.71	1.22	2.07	3.52	6.00	avg
$\rho = 0.00$	gap	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	time	45.2	38.8	38.6	29.5	29.3	28.4	27.4	26.3	26.4	25.9	31.6
$\rho = 0.35$	gap	0.3	0.2	0.3	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1
	time	48.0	47.6	49.4	44.1	39.3	30.7	29.0	29.1	27.3	28.0	37.3

Table 5: Optimality gap and computation time (in seconds) of sdp_2 with $n = 500$, $p = 100$, $s = k = 5$, $\lambda = \mu = 0$.

or near-optimal estimators, with an optimality gap at most 0.3%. In fact, with sdp_2 , we find that 97% of the estimators for $\rho = 0.00$ and 68% of the estimators with $\rho = 0.35$ are provably *optimal*¹¹ for (1). For a comparison, Hastie et al. (2017) report that, in their experiments, the MIO solver (with a time limit of three minutes) is able to prove optimality for only 35% of the instances generated with similar parameters. Although Hastie et al. (2017) do not report optimality gaps for the instances where optimality is not proven, we conjecture that such gaps are significantly larger than those reported in Table 5 due to weak relaxations with big- M formulations. In summary, for this class of instances, sdp_2 is able to produce optimal or practically optimal estimators of (1) in about 30 seconds.

5.3.4 RESULTS: ACCURACY METRICS

Figure 11 plots the relative risk, relative test error, the proportion of variance explained, and sparsity results as a function of the SNR for instances with $n = 500$, $p = 100$, $s = 5$ and $\rho = 0$. Figure 12 plots the same results for instances with $\rho = 0.35$. The setting with $\rho = 0.35$ was also used by Hastie et al. (2017).

We see that **elastic net** outperforms sdp_2 in low SNR settings, i.e., in SNR= 0.05 for $\rho = 0$ and SNR ≤ 0.14 for $\rho = 0.35$, but results in worse predictive performance for all other SNRs. Moreover, sdp_2 is able to recover the true sparsity pattern of β_0 for sufficiently large SNR, while **elastic net** is unable to do so. We also see that sdp_2 performs comparatively better than **elastic net** in instances with $\rho = 0$. Indeed, for large autocorrelations ρ , features where $(\beta_0)_i = 0$ still have predictive value thus the dense estimator obtained by **elastic net** retains a relatively good predictive performance (however, such dense solutions are undesirable from an *interpretability* perspective). In contrast, when $\rho = 0$, such features are simply noise, and **elastic net** results in overfitting, while methods that deliver sparse solutions such as sdp_2 perform much better in comparison. We also note that sdp_2 selects models corresponding to sparsities $k < s$ in low SNRs, while it consistently selects models with $k \approx s$ in high SNRs. We point out that, as suggested by Mazumder et al. (2022), the results for low SNR could potentially be improved by fitting models with $\mu > 0$.

Remark 14 *As an alternative to cross-validation, the sparsity of the statistical model can also be selected using information criteria such as AIC (Akaike, 1974) and BIC (Schwarz*

11. A solution is deemed optimal if $\text{gap} < 10^{-4}$, which is the default parameter in MIO solvers.




Figure 11: Relative risk, relative test error, proportion of variance explained and sparsity as a function of SNR, with $n = 500$, $p = 100$, $s = 5$ and $\rho = 0.00$.




Figure 12: Relative risk, relative test error, proportion of variance explained and sparsity as a function of SNR, with $n = 500$, $p = 100$, $s = 5$ and $\rho = 0.35$.

et al., 1978). As shown by Gómez and Prokopyev (2021), this can be accomplished efficiently by combining the relaxation sdp_r with fractional optimization techniques. We refer the reader to Gómez and Prokopyev (2021) for detailed computational and statistical results.

5.4 Summary and extensions

The conic formulations sdp_2 and sdp_{LB} are able to deliver near-optimal solutions to problem (1) with p in the hundreds, and results in substantially better performance than MIO methods (computational times and numerical stability) in instances with low or no ℓ_2 regularization. Such performance makes the approaches directly applicable in several high-stakes domains where interpretability is a major consideration (e.g., $p \approx 50$ in the highly-publicized COMPAS recidivism case Wexler (2017); $p \approx 20$ in the setting considered by Cozad et al. (2014), where the output of the regression is used to optimize a carbon capture adsorber). Due to the computational challenges with solving conic optimization problems via second order methods, the solution approach presented in this paper (using an off-the-shelf solver) does not scale to larger instances. Nonetheless, as we now point out, the formulations presented here may serve as a basis for methods that scale to larger values of p , or tackle statistical problems other than sparse regression.

First, tailored implementations that do not rely on off-the-shelf solvers are possible. For example, Liu et al. (2023) consider sparse inference problems with graphical models, which can be interpreted as a special case of (1) where matrix $\mathbf{X}^\top \mathbf{X}$ is sparse. They use sdp_2 as a base relaxation of their method and develop a tailored primal-dual method to solve it: the resulting approach scales comfortably to problems with p in the thousands. Similarly, Wei and Küçükyavuz (2023) use outer-approximations to implement the rank-one relaxations, which (coupled with state-of-the-art branch-and-bound algorithms) can solve mixed-integer quadratic optimization problems with up to 500 variables to optimality.

Second, sdp_r can be used as a subroutine of a more sophisticated method. For example, Hazimeh and Mazumder (2020) propose a method that approximately solves problem (1), scales to problems with $p \approx 10^5$ and produces “combinatorially local” solutions. A key idea of their approach is to solve MIO problems involving only a small subset of the variables to escape local minima. In a similar vein, it would be possible to use relaxations sdp_2 with a small subset of the variables to identify descent directions.

Finally, we point out that both the quadratic loss function and sparsity are used pervasively in machine learning; thus, the theory and methods developed in this paper may be generally applicable. For example, Bertsimas et al. (2021) incorporate some of the ideas discussed in §4 to develop algorithms that solve matrix completion problems to certifiable optimality.

6 Conclusions

In this paper we derive strong convex relaxations for sparse regression. The relaxations are based on the ideal formulations for rank-one quadratic terms with indicator variables. The new relaxations are formulated as semidefinite optimization problems in an extended space and are stronger and more general than the state-of-the-art formulations. In our computational experiments, the proposed conic formulations outperform the existing ap-

proaches, both in terms of accurately approximating the best subset selection problems and of achieving desirable estimation properties in statistical inference problems with sparsity.

Acknowledgments and Disclosure of Funding

A. Atamtürk is supported, in part, by NSF AI Institute for Advances in Optimization Award 2112533 and Grant N00014-24-1-2149 from the Office of Naval Research. A. Gómez is supported, in part, by Grants No. 1818700 and 2006762 from the National Science Foundation and grant No FA9550-24-1-0086 from the Air Force Office of Scientific Research.

References

- Sina Aghaei, Mohammad Javad Azizi, and Phebe Vayanos. Learning optimal and fair decision trees for non-discriminative decision-making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1418–1426, 2019.
- Hirotsugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–723, 1974.
- M Selim Aktürk, Alper Atamtürk, and Sinan Gürel. A strong conic quadratic reformulation for machine-job assignment with controllable processing times. *Operations Research Letters*, 37:187–191, 2009.
- Alper Atamtürk and Andrés Gómez. Strong formulations for quadratic optimization with M-matrices and indicator variables. *Mathematical Programming*, 170:141–176, 2018.
- Alper Atamtürk and Andrés Gómez. Safe screening rules for ℓ_0 -regression from perspective relaxations. In *International Conference on Machine Learning*, pages 421–430. PMLR 119, 2020.
- Alper Atamtürk and Vishnu Narayanan. Cuts for conic mixed integer programming. In M. Fischetti and D. P. Williamson, editors, *Proceedings of the 12th International IPCO Conference*, pages 16–29, 2007.
- Alper Atamtürk, Andrés Gómez, and Shaoning Han. Sparse and smooth signal estimation: Convexification of ℓ_0 -formulations. *Journal of Machine Learning Research*, 22(52):1–43, 2021.
- Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 15th International Conference, CPAIOR 2018, Delft, The Netherlands, June 26–29, 2018, Proceedings 15*, pages 35–51. Springer, 2018.
- Walid Ben-Ameur. Subset selection and the cone of factor-width-k matrices. *SIAM Journal on Optimization*, 34(1):817–843, 2024.

- Dimitris Bertsimas and Angela King. OR forum – an algorithmic approach to linear regression. *Operations Research*, 64:2–16, 2015.
- Dimitris Bertsimas and Bart Van Parys. Sparse high-dimensional regression. *The Annals of Statistics*, 48(1):300–323, 2020.
- Dimitris Bertsimas, Angela King, Rahul Mazumder, et al. Best subset selection via a modern optimization lens. *The Annals of Statistics*, 44:813–852, 2016.
- Dimitris Bertsimas, Ryan Cory-Wright, and Jean Pauphilet. Mixed-projection conic optimization: A new paradigm for modeling rank constraints. *Operations Research*, 70(6):3321–3344, 2021.
- Robert E Bixby. A brief history of linear and mixed-integer programming computation. *Documenta Mathematica*, pages 107–121, 2012.
- Valeriia Cherepanova, Roman Levin, Gowthami Somepalli, Jonas Geiping, C Bayan Bruss, Andrew G Wilson, Tom Goldstein, and Micah Goldblum. A performance-driven benchmark for feature selection in tabular deep learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Michaël Chichignoud, Johannes Lederer, and Martin J Wainwright. A practical scheme and fast algorithm to tune the lasso with optimality guarantees. *The Journal of Machine Learning Research*, 17:8162–8181, 2016.
- Alison Cozad, Nikolaos V Sahinidis, and David C Miller. Learning surrogate models for simulation-based optimization. *AIChE Journal*, 60:2211–2227, 2014.
- Anna Deza and Alper Atamtürk. Safe screening for logistic regression with $\ell_0 - \ell_2$ regularization. In *KDIR*, pages 119–126, 2022.
- Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Vu C Dinh and Lam S Ho. Consistent feature selection for analytic deep neural networks. *Advances in Neural Information Processing Systems*, 33:2420–2431, 2020.
- Hongbo Dong, Kun Chen, and Jeff Linderoth. Regularization vs. relaxation: A conic optimization perspective of statistical variable selection. *arXiv preprint arXiv:1510.06083*, 2015.
- Bradley Efron, Trevor Hastie, Iain Johnstone, Robert Tibshirani, et al. Least angle regression. *The Annals of Statistics*, 32:407–499, 2004.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2001.
- Antonio Frangioni and Claudio Gentile. Perspective cuts for a class of convex 0–1 mixed integer programs. *Mathematical Programming*, 106:225–236, 2006.

- Antonio Frangioni, Claudio Gentile, and James Hungerford. Decompositions of semidefinite matrices and the perspective reformulation of nonseparable quadratic programs. *Mathematics of Operations Research*, 45(1):15–33, 2020.
- Diko E Frank and Jerome H Friedman. A statistical view of some chemometrics regression tools. *Technometrics*, 35:109–135, 1993.
- Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- Andrés Gómez. Strong formulations for conic quadratic optimization with indicator variables. *Mathematical Programming*, 188(1):193–226, 2021.
- Andrés Gómez and Oleg A Prokopyev. A mixed-integer fractional optimization approach to best subset selection. *INFORMS Journal on Computing*, 33(2):551–565, 2021.
- Oktay Günlük and Jeff Linderoth. Perspective reformulations of mixed integer nonlinear programs with indicator variables. *Mathematical Programming*, 124:183–205, 2010.
- Shaoning Han and Andrés Gómez. Compact extended formulations for low-rank functions with indicator variables. *Mathematics of Operations Research*, 2024.
- Shaoning Han, Andrés Gómez, and Alper Atamtürk. The equivalence of optimal perspective formulation and Shor’s SDP for quadratic programs with indicator variables. *Operations Research Letters*, 50(2):195–198, 2022.
- Shaoning Han, Andrés Gómez, and Alper Atamtürk. 2×2 -convexifications for convex quadratic optimization with indicator variables. *Mathematical Programming*, 202(1):95–134, 2023.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 1. Springer series in statistics New York, NY, USA, 2001.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC press, 2015.
- Trevor Hastie, Robert Tibshirani, and Ryan J Tibshirani. Extended comparisons of best subset selection, forward stepwise selection, and the lasso. *arXiv preprint arXiv:1707.08692*, 2017.
- Hussein Hazimeh and Rahul Mazumder. Fast best subset selection: Coordinate descent and local combinatorial optimization algorithms. *Operations Research*, 68(5):1517–1537, 2020.
- Hussein Hazimeh, Rahul Mazumder, and Ali Saab. Sparse regression at scale: Branch-and-bound rooted in first-order optimization. *Mathematical Programming*, 196(1):347–388, 2022.

- Hussein Hazimeh, Rahul Mazumder, and Tim Nonet. L0learn: A scalable package for sparse learning using l0 regularization. *Journal of Machine Learning Research*, 24(205): 1–8, 2023a.
- Hussein Hazimeh, Rahul Mazumder, and Peter Radchenko. Grouped variable selection with discrete optimization: Computational and statistical perspectives. *The Annals of Statistics*, 51(1):1–32, 2023b.
- Mohamed Hebiri, Sara Van De Geer, et al. The smooth-lasso and other $\ell_1 + \ell_2$ -penalized methods. *Electronic Journal of Statistics*, 5:1184–1226, 2011.
- Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67, 1970.
- Jian Huang, Yuling Jiao, Yanyan Liu, and Xiliang Lu. A constructive approach to L0 penalized regression. *The Journal of Machine Learning Research*, 19:403–439, 2018.
- David R Hunter and Runze Li. Variable selection using MM algorithms. *Annals of Statistics*, 33:1617, 2005.
- Luca Insolia, Ana Kenney, Francesca Chiaromonte, and Giovanni Felici. Simultaneous feature selection and outlier detection with optimality guarantees. *Biometrics*, 78(4): 1592–1603, 2022.
- Hyemin Jeon, Jeff Linderoth, and Andrew Miller. Quadratic cone cutting surfaces for quadratic programs with on–off constraints. *Discrete Optimization*, 24:32–50, 2017.
- Jonathan Kelner, Frederic Koehler, Raghu Meka, and Dhruv Rohatgi. Feature adaptation for sparse linear regression. *Advances in Neural Information Processing Systems*, 36, 2024.
- Burak Kocuk, Santanu S Dey, and X Andy Sun. Strong socp relaxations for the optimal power flow problem. *Operations Research*, 64:1177–1196, 2016.
- Burak Kocuk, Santanu S Dey, and X Andy Sun. Matrix minor reformulation and SOCP-based spatial branch-and-cut method for the AC optimal power flow problem. *Mathematical Programming Computation*, 10:557–596, 2018.
- Simge Küçükyavuz, Ali Shojaie, Hasan Manzour, Linchuan Wei, and Hao-Hsiang Wu. Consistent second-order conic integer programming for learning bayesian networks. *Journal of Machine Learning Research*, 24(322):1–38, 2023.
- Ismael Lemhadri, Feng Ruan, Louis Abraham, and Robert Tibshirani. Lassonet: A neural network with feature sparsity. *Journal of Machine Learning Research*, 22(127):1–29, 2021.
- Xiaodong Lin, Minh Pham, and Andrzej Ruszczyński. Alternating linearization for structured regularization problems. *The Journal of Machine Learning Research*, 15:3447–3481, 2014.

- Peijing Liu, Salar Fattahi, Andrés Gómez, and Simge Küçükyavuz. A graph-based decomposition method for convex quadratic optimization with indicators. *Mathematical Programming*, 200(2):669–701, 2023.
- Michele Lombardi, Michela Milano, and Andrea Bartolini. Empirical decision model learning. *Artificial Intelligence*, 244:343–367, 2017.
- Donato Maragno, Holly Wiberg, Dimitris Bertsimas, S Ilker Birbil, Dick den Hertog, and Adejuyigbe Fajemisin. Mixed-integer optimization with constraint learning. *arXiv preprint arXiv:2111.04469*, 2021.
- Rahul Mazumder, Jerome H Friedman, and Trevor Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106:1125–1138, 2011.
- Rahul Mazumder, Peter Radchenko, and Antoine Dedieu. Subset selection with shrinkage: Sparse linear modeling when the SNR is low. *Operations Research*, 2022.
- Alan Miller. *Subset Selection in Regression*. CRC Press, 2002.
- Ryuhei Miyashiro and Yuichi Takano. Mixed integer second-order cone programming formulations for variable selection in linear regression. *European Journal of Operational Research*, 247:721–731, 2015.
- Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24:227–234, 1995.
- Daniel Nevo and Ya’acov Ritov. Identifying a minimal class of models for high-dimensional data. *The Journal of Machine Learning Research*, 18:797–825, 2017.
- Oscar Hernan Madrid Padilla, James Sharpnack, James G Scott, and Ryan J Tibshirani. The dfs fused lasso: Linear-time denoising over general graphs. *The Journal of Machine Learning Research*, 18:176–1, 2017.
- Mert Pilanci, Martin J Wainwright, and Laurent El Ghaoui. Sparse learning via boolean relaxations. *Mathematical Programming*, 151:63–87, 2015.
- Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Gideon Schwarz et al. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978.
- Soroosh Shafiee and Fatma Kılınç-Karzan. Constrained optimization of rank-one functions with indicator variables. *Mathematical Programming*, pages 1–47, 2024.
- Naum Z Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25:1–11, 1987.
- Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67:91–108, 2005.
- Ryan J Tibshirani and Jonathan Taylor. The solution path of the generalized Lasso. *The Annals of Statistics*, 39(3):1335–1371, 2011.
- Andreas M Tillmann, Daniel Bienstock, Andrea Lodi, and Alexandra Schwartz. Cardinality minimization, constraints, and regularization: a survey. *SIAM Review*, 66(3):403–477, 2024.
- Berk Ustun and Cynthia Rudin. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning*, 102:349–391, 2016.
- Berk Ustun and Cynthia Rudin. Learning optimized risk scores. *Journal of Machine Learning Research*, 20(150):1–75, 2019.
- Catherine Wanjiru, William Ogallo, Girmaw Abebe Tadesse, Charles Wachira, Isaiah Onando Mulang, and Aisha Walcott-Bryant. Automated supervised feature selection for differentiated patterns of care. *arXiv preprint arXiv:2111.03495*, 2021.
- Linchuan Wei and Simge Küçükyavuz. An outer approximation method for solving mixed-integer convex quadratic programs with indicators. *arXiv preprint arXiv:2312.04812*, 2023.
- Linchuan Wei, Andrés Gómez, and Simge Küçükyavuz. Ideal formulations for constrained convex optimization problems with indicator variables. *Mathematical Programming*, 192(1):57–88, 2022.
- Linchuan Wei, Alper Atamtürk, Andrés Gómez, and Simge Küçükyavuz. On the convex hull of convex quadratic optimization problems with indicators. *Mathematical Programming*, 204(1):703–737, 2024.
- Rebecca Wexler. When a computer program keeps you in jail. *New York Times*, 2017.
- Weijun Xie and Xinwei Deng. Scalable algorithms for the sparse ridge regression. *SIAM Journal on Optimization*, 30(4):3359–3386, 2020.
- Jiaming Zeng, Berk Ustun, and Cynthia Rudin. Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 180(3):689–722, 2017.
- Cun-Hui Zhang, Jian Huang, et al. The sparsity and bias of the Lasso selection in high-dimensional linear regression. *The Annals of Statistics*, 36:1567–1594, 2008.
- Cun-Hui Zhang, Tong Zhang, et al. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, 27:576–593, 2012.
- Cun-Hui Zhang et al. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38:894–942, 2010.

- Yuchen Zhang, Martin J Wainwright, and Michael I Jordan. Lower bounds on the performance of polynomial-time algorithms for sparse linear regression. In *Conference on Learning Theory*, pages 921–948, 2014.
- Peng Zhao and Bin Yu. On model selection consistency of Lasso. *The Journal of Machine Learning Research*, 7(Nov):2541–2563, 2006.
- Jin Zhu, Xueqin Wang, Liyuan Hu, Junhao Huang, Kangkang Jiang, Yanhang Zhang, Shiyun Lin, and Junxian Zhu. abess: a fast best-subset selection library in python and r. *Journal of Machine Learning Research*, 23(202):1–7, 2022.
- Junxian Zhu, Canhong Wen, Jin Zhu, Heping Zhang, and Xueqin Wang. A polynomial algorithm for best-subset selection problem. *Proceedings of the National Academy of Sciences*, 117(52):33117–33123, 2020.
- Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101:1418–1429, 2006.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67:301–320, 2005.
- Hui Zou and Runze Li. One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of Statistics*, 36:1509–1533, 2008.